S. HRG. 101-613

DEPARTMENT OF ENERGY HIGH-PERFORMANCE COMPUTING ACT OF 1989

HEARING

BEFORE THE

SUBCOMMITTEE ON
ENERGY RESEARCH AND DEVELOPMENT

OF THE

COMMITTEE ON
ENERGY AND NATURAL RESOURCES
UNITED STATES SENATE

ONE HUNDRED FIRST CONGRESS

SECOND SESSION

ON

S. 1976

TO PROVIDE FOR CONTINUED UNITED STATES LEADERSHIP IN HIGH-PERFORMANCE COMPUTING

MARCH 6, 1990

Printed for the use of the
Committee on Energy and Natural Resources

U.S. GOVERNMENT PRINTING OFFICE

28-910 WASHINGTON : 1990

For sale by the Superintendent of Documents, Congressional Sales Office
U.S. Government Printing Office, Washington, DC 20402

S311-28.

COMMITTEE ON ENERGY AND NATURAL RESOURCES

J. BENNETT JOHNSTON, Louisiana, *Chairman*

DALE BUMPERS, Arkansas	JAMES A. McCLURE, Idaho
WENDELL H. FORD, Kentucky	MARK O. HATFIELD, Oregon
HOWARD M. METZENBAUM, Ohio	PETE V. DOMENICI, New Mexico
BILL BRADLEY, New Jersey	MALCOLM WALLOP, Wyoming
JEFF BINGAMAN, New Mexico	FRANK H. MURKOWSKI, Alaska
TIMOTHY E. WIRTH, Colorado	DON NICKLES, Oklahoma
KENT CONRAD, North Dakota	CONRAD BURNS, Montana
HOWELL T. HEFLIN, Alabama	JAKE GARN, Utah
JOHN D. ROCKEFELLER IV, West Virginia	MITCH McCONNELL, Kentucky

DARYL OWEN, *Staff Director*
D. MICHAEL HARVEY, *Chief Counsel*
FRANK M. CUSHING, *Staff Director for the Minority*
GARY G. ELLSWORTH, *Chief Counsel for the Minority*

SUBCOMMITTEE ON ENERGY RESEARCH AND DEVELOPMENT

WENDELL H. FORD, Kentucky, *Chairman*
JOHN D. ROCKEFELLER IV, West Virginia, Vice Chairman

DALE BUMPERS, Arkansas	PETE V. DOMENICI, New Mexico
HOWARD M. METZENBAUM, Ohio	MITCH McCONNELL, Kentucky
TIMOTHY E. WIRTH, Colorado	DON NICKLES, Oklahoma
HOWELL T. HEFLIN, Alabama	CONRAD BURNS, Montana
	JAKE GARN, Utah

J. BENNETT JOHNSTON and JAMES A. MCCLURE are Ex Officio Members of the Subcommittee

PAUL BARNETT, *Counsel, Majority*
RICHARD D. GRUNDY, *Senior Professional Staff Member, Minority*

CONTENTS

Page

S. 1976 10

STATEMENTS

Bingaman, Hon. Jeff, U.S. Senator from New Mexico 5
Bromley, Dr. D. Allan, Director, Office of Science and Technology Policy 37
Connolly, Dr. John W.D., director, Center on Computational Sciences, University of Kentucky 130
Ford, Hon. Wendell H., U.S. Senator from Kentucky 1
Gore, Hon. Albert, U.S. Senator from Tennessee 27
Hecker, Dr. Siegfried S., Director, Los Alamos National Laboratory 106
Johnston, Hon. J. Bennett, U.S. Senator from Louisiana 2
McClure, Hon. James A., U.S. Senator from Idaho 8
Nelson, Dr. David B., Executive Director, Office of Energy Research, Department of Energy 55
Paluck, Robert J., chairman, president and chief executive officer, Convex Computer Corp 90
Rollwagen, John A., chairman and chief executive officer, Cray Research, Inc 73

APPENDIXES

APPENDIX I

Responses to additional questions 149

APPENDIX II

Additional material submitted for the record 179

DEPARTMENT OF ENERGY HIGH-PERFORMANCE COMPUTING ACT OF 1989

TUESDAY, MARCH 6, 1990

U.S. SENATE,
SUBCOMMITTEE ON ENERGY RESEARCH AND DEVELOPMENT
COMMITTEE ON ENERGY AND NATURAL RESOURCES,
Washington, D.C.

The subcommittee met, pursuant to notice, at 2:05 p.m., in room SD-366, Dirksen Senate Office Building, Hon. Wendell H. Ford, presiding.

OPENING STATEMENT OF HON. WENDELL H. FORD, U.S. SENATOR FROM KENTUCKY

Senator FORD. Good afternoon ladies and gentlemen. We have quite a group of witnesses here today. I was sorry to learn that one of our witnesses could not make it, Sid Karen, who heads up the supercomputer center in San Diego. He received a very serious injury last week, breaking his leg, and we all wish him well. I hope he will send us his comments for the record and we intend to pursue that request.

I am delighted to have the distinguished Senator from Tennessee, Senator Gore, here. I know he has worked hard on this issue and that he and I sit on the Commerce Committee together. He has held quite a few hearings on the Commerce Committee, and I think it has made all of our jobs a lot easier because of his efforts.

So, let me ask, not only Senator Gore, but the other witnesses, since there is a possibility of votes shortly after 2:15 for the rest of the afternoon, that the witnesses give a short summary of their written testimony, give us time to ask a few questions, and then we can proceed.

Senator Gore, we are delighted to have you, and if you want to proceed with your statement, I would be very pleased to receive it.

[Prepared statements of Senators Johnston, Bingaman, and McClure and the text of S. 1976 follow:]

STATEMENT OF SENATOR J. BENNETT JOHNSTON

Subcommittee on Energy Research and Development

Hearing on S. 1976, the "Department of Energy High-Performance Computing Act of 1989"

GOOD AFTERNOON. I AM HONORED TO BE PRESENT TO HEAR THE DISTINGUISHED GUESTS FEATURED TODAY. I VERY MUCH LOOK FORWARD TO HEARING THEIR THOUGHTS ON MY BILL. THE WITNESSES WHO HAVE COME HERE TO TESTIFY SPEAK TO THE IMPORTANCE OF THIS ISSUE. THE LEVEL OF ENERGY AND ENTHUSIASM THIS ISSUE HAS GENERATED IS ASTOUNDING.

THE SITUATION CONGRESS FACES TODAY IS MUCH LIKE THAT THE 73RD CONGRESS FACED BEFORE ADOPTING THE TELECOMMUNICATIONS ACT OF 1934, WHICH CREATED THE NATION'S TELECOMMUNICATIONS SYSTEM.

THOUGH I WAS NOT HERE WHEN IT PASSED, THE COMMUNICATIONS ACT OF 1934 WAS THE LAST MAJOR TELECOMMUNICATIONS POLICY ACT ADOPTED BY CONGRESS. THE 73RD CONGRESS CREATED A POLICY THAT SPAWNED THE BEST TELECOMMUNICATIONS SYSTEM IN THE WORLD. THAT POLICY - ENCOURAGING RESEARCH AND DEVELOPMENT, AND, ENCOURAGING CAPITAL INVESTMENT - FUELED THE DEVELOPMENT OF A NETWORK THAT IS THE MOST POWERFUL COMPUTER IN THE WORLD. THE TELEPHONE SYSTEM, IS AFTER ALL, A GIGANTIC COMPUTER. IT IS AN UNBEATABLE COMBINATION - AMERICAN ACADEMIA, GOVERNMENT, AND INDUSTRY ALL WORKING TOGETHER FOR THE SAME CAUSE.

THE PURPOSE OF MY BILL IS NOT TO HAVE THE U.S. GOVERNMENT BEAR THE MAJOR INVESTMENT BURDEN. THE PURPOSE IS TO ACT AS A CATALYST: TO REDUCE RESEARCH AND DEVELOPMENT RISK TO AMERICAN

INDUSTRY; TO SUPPORT JOINT VENTURES BETWEEN INDUSTRY, GOVERNMENT AND ACADEMIA; TO HAVE A BALANCED RESEARCH AND DEVELOPMENT PROGRAM IN NEW COMPUTING TECHNOLOGIES; TO HELP EDUCATE A NEW GENERATION OF COMPUTER AND ENGINEERING SCIENTISTS; TO ESTABLISH A NATIONAL RESEARCH AND EDUCATION NETWORK. THE RETURN ON OUR INVESTMENT WILL BE ENORMOUS. FAILURE TO INVEST WILL COST US DEARLY.

THE FEDERAL GOVERNMENT RELIES ON SUPERCOMPUTING FOR DEFENSE AND SECURITY. NO DOUBT, THE SUPERCOMPUTER IS AS IMPORTANT TO NATIONAL SECURITY AS ARE BATTLESHIPS, SUBMARINES AND AIRCRAFT. FOR THAT REASON ALONE IT IS IMPERATIVE THAT THE U.S. STAY AT THE FOREFRONT OF SUPERCOMPUTING TECHNOLOGY. SUPERCOMPUTING IS EQUALLY AS IMPORTANT TO OUR NATION'S SCIENTIFIC COMMUNITY.

IT IS ONLY BECAUSE OF SUPERCOMPUTING THAT WE CAN ACCURATELY STUDY THE PHENOMENON OF GLOBAL WARMING. THE SUPERCONDUCTING SUPERCOLLIDER CANNOT WORK WITHOUT SUPERCOMPUTERS. INDUSTRY TOO HAS RECOGNIZED THIS TOOL'S POWER, AND IS BEGINNING TO MAKE WIDESPREAD USE OF IT.

WITH PARTICULAR INTEREST, I NOTICED THAT THE OIL INDUSTRY IS ONE OF THE BIGGEST USERS OF SUPERCOMPUTERS. SUPERCOMPUTERS HAVE PAID FOR THEMSELVES MANY TIMES OVER IN THIS INDUSTRY. USING SUPERCOMPUTERS, OIL FIELD PRODUCTION HAS INCREASED AND, COSTLY DRY HOLES HAVE BEEN AVOIDED.

NO ONE QUESTIONS THE IMPORTANCE OF SUPERCOMPUTING. SO, IT WAS WITH ALARM, THAT I LEARNED WHEN CONTROL DATA CLOSED ITS SUPERCOMPUTER OPERATIONS ONLY ONE SUPERCOMPUTER MANUFACTURER REMAINS. CONTROL DATA INVESTED HUNDREDS OF MILLIONS OF DOLLARS

IN AN ENTERPRISE THAT SHOULD HAVE WORKED. IT WAS AN INVESTMENT WE SHOULD NOT HAVE LOST.

CRAY RESEARCH, THE FOUNDER OF THE SUPERCOMPUTER, IS THAT ONE REMAINING SUPERCOMPUTER COMPANY. WHEN THE 1980'S STARTED, CRAY DOMINATED THE WORLD MARKET, WITH NO ONE IN SIGHT. A LITTLE OVER FOUR YEARS AGO THE JAPANESE ENTERED THE FIELD. TODAY, SOME ESTIMATE WE HAVE BUT A TWO YEAR LEAD. THAT LEAD CANNOT COME ANY CLOSER IF INDUSTRY IS TO REMAIN COMPETITIVE. THE CLOSER OTHERS COME, THE MORE OUR NATIONAL SECURITY IS THREATENED.

IT IS THE PURPOSE OF MY BILL TO NOT ONLY STOP THE NARROWING OF THAT GAP, BUT TO WIDEN IT. CRAY ACHIEVED ITS PREEMINENT STATUS BY WORKING WITH THE DEPARTMENT OF ENERGY TO DEVELOP THE WORLD'S FIRST SUPERCOMPUTER. THE DEPARTMENT HAS CONTINUED TO SUPPORT SUPERCOMPUTING AS HAVE OTHER FEDERAL AGENCIES. BUT THE AMOUNT OR KIND OF HELP BEING GIVEN TODAY IS NO LONGER ENOUGH. HOPEFULLY, TODAY, WE CAN BETTER LEARN WHAT KIND OF HELP CONGRESS CAN AND SHOULD GIVE, SPECIFICALLY THE DEPARTMENT OF ENERGY, AND HOW FAR MY BILL GOES TOWARD GIVING THAT HELP.

JUST AS THE 73RD CONGRESS SEIZED THE OPPORTUNITY TO SPAWN THE WORLD'S GREATEST TELECOMMUNICATION'S NETWORK, THE 101ST CONGRESS NOW HAS THE OPPORTUNITY TO SPAWN THE WORLD'S GRANDEST TECHNOLOGICAL EFFORT.

Statement by Senator Jeff Bingaman (D-NM)

The Department of Energy High-Performance Computing Act
March 6, 1990

I would like to thank the Chairman for scheduling this hearing.

The legislation before us today, the Department of Energy High-Performance Computing Act - S.1976, is an important step in addressing the Department of Energy's role in the development and use of high-performance computing. I commend Senator Johnston for introducing this legislation.

S. 1976 is part of a larger effort, spearheaded by Senator Gore, to create a comprehensive government-wide policy for the development of supercomputers. I am pleased to be a cosponsor of Senator Gore's bill, which effectively implements the plan on high-performance computing developed by the Office of Science and Technology Policy. I am concerned, however, that this plan is not adequately reflected in the President's budget submission

I have an overlapping set of interests in this and Senator Gore's legislation. As a member of the Energy Committee and with two premier Department of Energy laboratories, Los Alamos and Sandia, in my home state of New

Mexico, I have a strong interest in further strengthening the labs' role in the development and use of high-performance computers.

I would like to point out that later in the hearing we be receiving testimony from the Director of Los Alamos National Laboratory, Dr. Siegfried Hecker. Los Alamos has been important in fostering the development of supercomputers, largely through their weapons systems' work but increasingly through work in civilian areas such as the human genome project. I look forward to Dr. Hecker's testimony.

I also have a strong interest in this issue through my roles as Chairman of the Defense Industry and Technology subcommittee of the Armed Services Committee, which has jurisdiction over the DoD research component of the high-performance computing plan, and as Chairman of the Government Information and Regulation subcommittee of the Governmental Affairs Committee, which has jurisdiction over information technology policy for the Federal government.

I believe it is important that we produce a overall policy for high-performance computing. I hope to be able to address the DoD component of this policy in this year's Defense bill and promise to work closely with you, Mr.

Chairman, and with Senators Johnston and Gore to bring such a policy about.

I thank you for your courtesy and again commend you for holding this hearing.

Statement of
Senator James A. McClure (R-ID)

Hearing on S. 1976, the Department of Energy
High-Performance Computing Act of 1989

March 6, 1990

Mr. Chairman, I am pleased the Committee is holding today's hearing on S. 1976, the Department of Energy High-Performance Computing Act.

In an era of increasing global competition, it is essential that full use be made of available Federal high-speed computing resources and that appropriate action be taken to promote the development and dissemination of new computing technologies.

The role of the Federal government to date has largely been that of a consumer; technological advancements have been made to meet the Government's needs, and the private sector has benefitted from the spin-offs -- but the Government has not set priorities or directions. In a nutshell, that is what this legislation is all about.

This legislation proposes to take the Federal government a step further than where it is today; to make it more active in setting directions and priorities. While this has some appeal, it also raises a number of

complex policy issues regarding the Government's proper role and responsibilities vis-a-vis the private sector, as well as the most effective and efficient means of accomplishing that objective.

I look forward to hearing our witnesses' views on these issues, and receiving their suggestions as to how this legislation may be improved.

II

101ST CONGRESS
1ST SESSION

S. 1976

To provide for continued United States leadership in high-performance computing.

IN THE SENATE OF THE UNITED STATES

NOVEMBER 21 (legislative day, NOVEMBER 6), 1989

Mr. JOHNSTON (for himself, Mr. GORE, and Mr. MCCLURE) introduced the following bill; which was read twice and referred to the Committee on Energy and Natural Resources

---

A BILL

To provide for continued United States leadership in high-performance computing.

Be it enacted by the Senate and House of Representatives of the United States of America in Congress assembled, That this Act may be referred to as the "Department of Energy High-Performance Computing Act of 1989".

FINDINGS AND PURPOSES

SEC. 2. (a) The Congress finds and declares the following:

(1) In the last twenty years, computing technology has transformed America's research laboratories,

factories, and offices, and become indispensable to our way of life.

(2) Rapid advances in computing technology have resulted in uses for computers unimaginable only five or ten years ago. Many of these advances are a result of research and development on supercomputers, advanced computer software, and other aspects of high-performance computing technology.

(3) High-performance computing is a powerful tool to increase productivity in industrial design and manufacturing, scientific research, communications, and information management.

(4) The United States currently leads the world in the development and use of high-performance computing. However, that lead is increasingly being challenged, and American firms share of the multi- billion-dollar world market for both high-performance computer systems and other computers is shrinking.

(5) In order to strengthen America's computer industry and to assist the entire manufacturing sector, the Federal Government must provide leadership in the development and application of high-performance computer technology. In particular, the Federal Government should create a National High-Performance Computing Program to support the development of a high-

capacity, national research and education computer network; facilitate the development of software for research, education, and industrial applications; continue to fund basic research; and provide for the training of computer scientists and computational scientists.

(6) Several Federal agencies have ongoing high-performance computing research and development programs which can contribute to a National High-Performance Computing Program. Such a program would provide additional funding for these existing programs, create new research and development programs, and improve coordination between the various agency programs.

(7) A September 1989 report by the Office of Science and Technology Policy entitled "The Federal High Performance Computing Program" outlining a research and development plan provides a framework for such a program.

(8) The Department of Energy, in order to fulfill its mission to conduct energy research and direct the Nation's nuclear weapons program, has established several high-performance computing research and development programs. High-energy physics, materials sciences, fusion energy research, human genetics research, oil and gas exploration, nuclear reactor design,

and nuclear weapons design all rely heavily on high-performance computing.

(9) The Department of Energy has extensive high-performance computing facilities and has played a key role in developing software and applications for supercomputers. It has funded research in mathematical and computational sciences, has developed new designs for supercomputers, and has established advanced computer networks for connecting supercomputers and other computers throughout the country.

(10) By building upon existing Department of Energy high-performance computing research and development programs, the Department of Energy can play a key role in a National High-Performance Computing Program.

(b) PURPOSE.—It is the purpose of the Congress in this Act to authorize the Secretary of Energy to—

(1) develop a long-range strategy for research, development, and application of high-performance computing;

(2) implement that strategy in conjunction with other Federal agencies as part of a National High-Performance Computing Program; and

(3) ensure the appropriate transfer of high-performance computing technology to United States industry.

(c) DEFINITIONS.—

(1) "Secretary" means the Secretary of Energy.

(2) "Department" means the Department of Energy.

DEPARTMENT OF HIGH-PERFORMANCE COMPUTING PLAN

SEC. 3. (a) The Secretary is authorized to immediately develop and implement long-range plan for high-performance computing at the Department of Energy. The Secretary shall develop the plan within one year after the date of enactment of this Act. The plan shall cover the fiscal year the plan is implemented and at least the next four years. The plan shall thereafter be updated annually.

(b) The plan shall—

(1) summarize ongoing high-performance computing programs at the Department of Energy;

(2) detail the Department of Energy's contribution to a National High-Performance Computing Program to expand Federal support for research, development, and application of high-performance computing technology in order to—

(A) establish a high-capacity national research and education computer network;

(B) develop data bases, services, and research facilities which would be available for access over such a national network;

(C) stimulate research on software technology;

(D) promote the more rapid development and wider distribution of computer software;

(E) accelerate the development of computer systems; and

(F) invest in basis research and education.

(3) establish the goals and priorities for research, development, and application of high-performance computing at the Department of Energy for the time period covered by the plan;

(4) describe the levels of funding for each aspect of high-performance computing, including basic research, hardware and software development, education, acquisition and operating expenses for computers and computer networks, and education;

(5) define the role of each of the Department of Energy's national laboratories involved in research, development, and application of high-performance computing technology; and

(6) set a timetable for creation and implementation of technology transfer mechanisms to ensure that

the results of research funded under the plan are readily available to United States industry.

THE NATIONAL RESEARCH AND EDUCATION NETWORK

SEC. 4. (a) As part of a National High-Performance Computing Program, the Secretary shall, in cooperation with the Director of the National Science Foundation, the Secretary of Defense, the Secretary of the Department of Commerce, the Administrator of the National Aeronautics and Space Administration, and other appropriate agencies, provide for the establishment of a national multi-gigabit-per-second research and education computer network by 1996, to be known as the National Research and Education Network (NREN).

(b) The network shall—

(1) link government, industry, and the higher education community;

(2) provide computer users at more than one thousand universities, Federal laboratories, and industrial laboratories with access to supercomputers, computer data bases, and other research facilities;

(3) be developed in close cooperation with the computer and telecommunications industry;

(4) be designed and developed with the advice of potential users in government, industry, and the higher education community;

(5) be established in a manner which fosters and maintains competition in high speed data networking within the telecommunications industry;

(6) have accounting mechanisms which allow users or groups of users to be charged for their usage of the network, where appropriate; and

(7) be phased out when commercial networks can meet the networking needs of American researchers.

(c) The Department of Energy shall—

(1) provide networking support for the energy research community;

(2) provide for interconnection of existing computer networks run by the Department and other agencies, where appropriate;

(3) participate, with other Federal agencies, in the development and testing of advanced prototype networks;

(4) conduct research and development of advanced networking technology, particularly for supercomputers;

(5) develop technology to support computer-based collaboration that allows researchers around the Nation to share information and instrumentation using computer networks; and

(6) take an active role in the interagency coordinating committees established to develop the National Research and Education Network.

(d) There is authorized to be appropriated to the Secretary for the purposes of this title, $10,000,000 for fiscal year 1991, $15,000,000 for fiscal year 1992, $20,000,000 for fiscal year 1993, and $25,000,000 for fiscal year 1994, and $30,000,000 for fiscal year 1995.

SOFTWARE

SEC. 5. (a) In accordance with the plan outlined in section 3, the Secretary shall provide for research and development of high-performance computer software for application in high-energy physics, fusion energy research, engineering, materials sciences, astrophysics, climate modeling, genetics, and other fields. The Secretary shall also provide for the development of improved software tools and components to facilitate the development of software for high-performance computer systems.

(b) The Secretary shall define and provide advanced software technology support to research groups collaborating to address so-called Grand Challenge problems in science and engineering. A Grand Challenge is a fundamental problem in science and engineering, with broad economic and scientific impact, whose solution will require the application of the high-performance computing resources.

(1) The Grand Challenges to be addressed include by are not limited to—

(A) PREDICTION OF GLOBAL CHANGE.—The goal is to understand the coupled atmosphere, ocean, biosphere system in enough detail to be able to make long-range predictions about its behavior and determine its response to man-caused releases of carbon dioxide, methane, chlorofluorocarbons, and other gases.

(B) MATERIALS SCIENCES.—The goal is to use high-performance computing technology to improve our understanding of the atomic nature of materials, enabling the design and production of improved semiconductors, superconductors, ceramics, and other materials.

(C) HUMAN GENOME.—The goal is to use high-performance computing technology to analyze, store, and disseminate data on the molecular structure of the DNA that expresses the tens of thousands of genes carried by each human being. Identification of these genes would lead to a better understanding and possibly treatment of genetic diseases, cancer, and other diseases.

(D) NUCLEAR FUSION.—The goal is to use supercomputer models to understand the physics

of plasmas at the very high temperature required for nuclear fusion.

(E) ENERGY EFFICIENCY.—The goal is to use supercomputer models to better understand combustion in order to design more efficient engines and furnaces.

(F) ENHANCED OIL AND GAS RECOVERY.—The goal is to use supercomputer models to locate and better exploit oil and gas fields.

(2) The Secretary shall focus research on those Grand Challenges that are of greatest importance to the Nation, will benefit most from the application of high-performance computing, and are most consistent with the mission of the Department of Energy.

(3) The Secretary shall establish collaborative research groups consisting of scientists and engineers concerned with a particular Grand Challenge, software and systems engineers, and algorithm designers, and provide them with—

(A) computational and experimental facilities, including supercomputers for numerical modeling;

(B) access to the National Research and Educational Network and other computer networks, and;

(C) access to and technology for effectively utilizing scientific data bases.

(d) The Secretary shall establish programs to develop software tools and components to accelerate development of software for computers, especially supercomputers. Such programs would fund research on fundamental algorithms, models of computation, program analysis, and new programming languages. Particular emphasis should be given to development of programming languages, compilers, operating systems, and software tools for parallel computer systems.

(e) The Secretary shall establish high-performance computing research centers to accelerate the development and application of new generations of high-performance computing technology by enabling researchers to explore applications of this new technology.

(1) Most of these centers would be located within existing computer research organizations funded by the Department of Energy;

(2) These centers could facilitate research on the Grand Challenges and other applications;

(3) Both new and existing Department of Energy supercomputers centers shall help provide the national research community with access to supercomputers because researchers developing algorithms, software

tools, and operating systems require access to new generation technology; and

(4) These centers shall provide access to a variety of different high-performance computer systems with different computer architectures.

(f) There is authorized to be appropriated to the Secretary for research and development on scientific Grand Challenges, development of advanced software technology, and creation of high-performance computing research centers, in accordance with the purposes of this section, $30,000,000 for fiscal year 1991, $45,000,000 for fiscal year 1992, $60,000,000 for fiscal year 1993, $75,000,000 for fiscal year 1994, $90,000,000 for fiscal year 1995.

HIGH-PERFORMANCE COMPUTER SYSTEMS

SEC. 6. (a) In accordance with the plan required in section 3, the Secretary shall provide for support of research and development of high-performance computer systems. Funding shall be provided for—

(1) research and development in the national laboratories, universities, and industry on all aspects of high-performance computer systems including processors, memory, mass storage devices, input/output devices, and associated system software;

(2) increased research in—

(A) computer science,

(B) parallel computer architectures,

(C) optoelectronics, and

(D) mass storage technology; and

(3) development of tools for the rapid design, prototyping, and integration of high-performance computing systems.

(b) In addition, the Department of Energy shall purchase early market and production model computer systems and subsystems for use both in high-performance computing research centers and for other research programs within the Department. Such purchases will—

(1) stimulate hardware and software development by reducing the research and development risk of United States manufacturers developing high-performance computer systems;

(2) provide manufacturers with valuable tests of their new systems.

(c) There is authorized to be appropriated to the Secretary for research and development, procuring, and testing of high-performance computer systems, $15,000,000 for fiscal year 1991, $25,000,000 for fiscal year 1992, $35,000,000 for fiscal year 1993, $45,000,000 for fiscal year 1994, and $55,000,000 for fiscal year 1995.

BASIC RESEARCH AND EDUCATION

SEC. 7. (a) In order to address the long-term national need for high-performance computing, the Secretary shall—

(1) support basic research on computer technology, including research on advanced semiconductor computer chip designs, new materials for integrated circuits, improved integrated circuit fabrication techniques, photonics, and superconducting computer components;

(2) support basic research on computing technology, including basic research on algorithms, software languages and tools, architectures, systems software, networks, distributed computing, and symbolic processing;

(3) create technology transfer mechanisms to ensure that the results of basic research are readily available to United States industry;

(4) promote basic research in computer science, computational science, electrical engineering, and materials sciences; and

(5) educate and train more researchers in computer science and computational science by—

(A) making the national laboratories available to senior graduate students, postdoctoral fellows, and faculty from the Nation's universities;

(B) expanding summer science programs for high school students;

(C) providing computer facilities to universities throughout the country; and

(D) establishing more cooperative research programs with the academic computational science community.

(b) There is authorized to be appropriated to the Secretary for the purposes of this section, $10,000,000 for fiscal year 1991, $15,000,000 for fiscal year 1992, $20,000,000 for fiscal year 1993, $25,000,000 for fiscal year 1994, and $30,000,000 for fiscal year 1995.

GOVERNMENT AND PRIVATE SECTOR COOPERATION

SEC. 8. (a) The Secretary may cooperate with, solicit help from, provide funds to, or enter into contracts with private contractors, industry, government, universities, or any other person or entity the Secretary deems necessary in carrying out the Provisions of this Act.

(b) The Secretary shall cooperate with other Federal agencies in carrying out the provisions of this Act, particularly the National Science Foundation, the Department of Commerce (particularly the National Institute of Standards and Technology), the National Aeronautics and Space Administration, the Department of Defense, and the Office of Science and Technology Policy.

REPORT REQUIREMENT

SEC. 9. The Secretary shall within one year after the date of enactment of this Act, report to the Congress regard-

ing the implementation of this Act, and thereafter, provide annual reports to the Congress.

STATEMENT OF HON. ALBERT GORE, JR., U.S. SENATOR FROM TENNESSEE

Senator GORE: Well, thank you very much, Mr. Chairman, and I am awfully glad to be here with you today, and to have a good friend and colleague in the chair, and one who is familiar with this legislation, because of your work on the Commerce Committee, and I want to express my appreciation to the other members of this Committee for considering this matter expeditiously.

I would like to express my thanks to the staff of the Committee for working very hard and in such a cooperative manner on this. It really makes a tremendous difference. When you have legislation that has to come out of three different committees, it does make it difficult, but when everybody works together, it really makes a tremendous difference, and the suggestions and initiatives of this Committee have been extremely helpful.

Supercomputing and high speed networking are critical to America's future competitiveness. Industries like the aerospace industry, the automotive industry, the petroleum industry, pharmaceuticals, chemicals, all of them, and others will use supercomputing.

Here is the essence of the problem, Mr. Chairman. We are on the seventh or eighth generation supercomputer, and we are still on the first generation network. We cannot send the volumes of information we need to use over the information highways that are presently available.

Now, the leading supercomputers today are in the over $10 million, $15 million, some of them even more, but that is the range we are talking about. That kind of machine just five years from now is going to cost $300,000 or $400,000. They are going to be commonplace in lots of companies. Medium-size companies in Kentucky and Louisiana and Tennessee are going to have these machines because their competitors are going to have these machines. And, if their competitors are in Japan, then they are going to be able to talk to their cadcam designer, the people who are designing their processes, their suppliers, their materials people, they are going to be able to communicate with very high flows of information throughout their manufacturing base in Japan. Same in parts of Europe.

The United States has an opportunity with this legislation and with the parallel proposals that are coming out of the executive branch. They are just not funded yet. We will have a chance to leapfrog what Japan and Europe are trying to do.

I call it information super highways. When the interstate system was designed and built, back in the 50s and 60s, some people said, well, we will not have as much car traffic and truck traffic to justify building the interstates. But, of course, we did and the people who put that system in place realized that when it was in place, the use would grow rapidly, and of course now, the interstates are being expanded in so many areas. Same thing is true with this legislation.

Now, a couple of housekeeping measures. The bill before you today, S. 1976, is a companion to S.1067, and it encompasses a five year program to ensure that the United States does not forfeit its lead in supercomputing. It coordinates several different agencies

that have a part of this problem to solve, and a part of the opportunity to seize.

The National Science Foundation, NASA, the Department of Energy and the Department of Defense, particularly, DARPA, the Defense Advance Research Program's Agency, others, including NOAA, the USGS, National Institutes of Health, Library of Congress, would also have important roles to play. So, clearly, a multiagency program does have to be carefully coordinated, and that is why title 1 of the bill calls for the Office of Science and Technology Policy to plan and coordinate this program.

Luckily, your next witness is Dr. Allan Bromley. I have come to know him well, as you have, in his many appearances before the Commerce Committee, and he is, in my opinion, and I have told him this before, doing an outstanding job.

Senator FORD. He may be a good soldier too, if OMB wants to oppose this bill.

Senator GORE. Well, I urge you not to take it too seriously when he says what OMB wants him to say on this.

Senator FORD. It is very difficult though, as you know, to get his personal opinion to override OMB's opinion.

Senator GORE. Well, I tell you, Mr. Chairman, Dr. Bromley is developing into the Everett Koop of the Bush Administration. He has the requisite candor and expertise and courage, but Mr. Chairman, do not put him to the test today, because it is early in his tenure. He is still gaining momentum and strength, and we all know what his personal opinions are on this legislation.

Senator FORD. I think we just put him on the spot.

Senator GORE. Well, when he testified before the Commerce Committee, he said this should be America's number one priority in science and technology.

Every scientific association, every scientific endeavor in this country, will benefit from this legislation. Industries using large computers will benefit from this legislation. Right now, a group using a supercomputer, particularly if they are using graphics, or moving graphical displays, and that is the way you get information in a form that you and I can understand it, because these rows of numbers are useless to people. Machines can deal with rows of numbers, but the only way I can understand them is if they get them into a picture that makes sense. And, if that picture moves over time, it contains even more information. Well, that is one of the things supercomputers are really good at.

Let's say that a team doing that in one area is working on the same problem or set of problems as another team in another area, and they come up with a new wrinkle that they need to consult on. What do they do? Well, they have to somehow get that computer program into a magnetic disk or some kind of storage medium, and take it with themselves and get on an airplane and fly to wherever the other group is, and then crank it up on the other supercomputer. We can link those teams together and streamline their work, and the work of American industry.

We can also use it for education, Mr. Chairman. I want to see a school child in Owensboro, Kentucky, come home after school and be able to sit down in front of a device that is not much different from, and not much more expensive, than one of these Nintendo

machines available today, and have access to the entire Library of Congress and other data bases that, in a way, allows that child to learn at his or her own speed, to explore ideas and concepts and areas that come out of their own natural curiosity and desire to learn. That can be done. We know how to do it today. But we do not have the network in place that makes it possible.

This legislation will make it possible, and again, Dr. Bromley has been extremely helpful on this. Several years ago, we passed a bill out of the Commerce Committee that authorized a multi-agency study of this whole deal, and they did an excellent job on that, and they came back with a report that was extremely helpful and the whole idea has picked up speed since then.

Now, I am going to put the rest of my full statement in the record, but I want to again, compliment Senator Johnston and Senator McClure, who joined me in introducing S. 1976, which deals specifically with the Energy Department's part of this whole initiative. I am very proud that the Senate committees are teaming up, just as Federal agencies work together on programs of priority.

Now, DOE has a particularly important role, because for years, of course, DOE has been at the forefront in using supercomputers. They had the first CRAY, they now have dozens of supercomputers in use helping DOE scientists in a variety of different fields, which I spell out in some detail in the prepared statement.

At a hearing on S. 1067 last July, Dr. Karl Heinz Winkler, who, until recently, was at Las Alamos with DOE, demonstrated some of the technology that he and others have developed that will be used on this network.

A lot of the technology being developed in the Department of Energy in this particular area is not classified and is already being put to good use by American firms. I will give you many examples in the prepared statement of exactly that phenomenon. But in order to save you time, let me close out my prepared statement by simply saying that I really believe that this legislation and this initiative can do as much as any single step to assist American industry and educational institutions to move forward quickly in taking advantage of what computers can do for us in making America more competitive in helping us to educate our children. And, I appreciate your time and attention, Mr. Chairman. Thank you.

[The prepared statement of Senator Gore follows:]

28-910 O - 90 - 2

TESTIMONY BY SENATOR ALBERT GORE, JR. ON

S. 1976, THE DEPARTMENT OF ENERGY HIGH-PERFORMANCE COMPUTING ACT OF 1989

BEFORE THE SENATE SUBCOMMITTEE ON ENERGY RESEARCH AND DEVELOPMENT

TUESDAY, MARCH 6, 1990

I am extremely pleased to have this chance to testify on supercomputing before the members of the Energy Research and Development Subcommittee. I am especially glad to be able to testify before Senator Ford, a good friend and colleague who serves with me on the Commerce Committee. I know that he too is eager to see America stay at the leading edge in computing technology.

I am here today to briefly summarize S. 1976, the DOE High-performance Computing Act of 1989, and to tell you why supercomputing is so important to scientists, engineers, educators, businessmen -- indeed, to all Americans. The technology that we are discussing here today is critical not just to America's computer industry; it is a key component of any strategy to keep American companies competitive in world markets. Supercomputing and high-speed networking will help our aerospace industry, our automotive industry, our petroleum industry, our pharmaceutical industry, and dozens of other industries, both high-tech and low-tech.

The bill before us today, S. 1976, is a companion bill to S. 1067, the National High-Performance Computer Technology Act, which I introduced last May. That bill laid out an

ambitious, five-year program to ensure that the United States does not forfeit its lead in supercomputing. The program would involve the National Science Foundation, NASA, the Department of Energy, and the Department of Defense (especially the Defense Advanced Research Programs Agency, DARPA, for short). Other agencies, including NOAA, the United States Geological Survey, the National Institutes of Health, and the Library of Congress would also have important roles to play.

Clearly, a multi-agency program like this will need to be carefully coordinated. Hence, Title I of S. 1067 call for the White House Office of Science and Technology Policy to plan and coordinate this program. As you know, the Director of OSTP is Dr. Allan Bromley, who is your next witness and who, in just a few months on the job, has provided real leadership in high-performance computing and a wide range of other science and technology issues. For more than four years, OSTP has been working on a Federal high-performance computing program, and we have worked closely with Dr. Bromley and OSTP and other federal agencies as we drafted S. 1067 and S. 1976. I look forward to continuing to work with Dr. Bromley and OSTP to make this program a reality.

S. 1067 called for roughly doubling the amount of Federal funding for high-performance computing over the next five years. It authorizes a total of $1.75 billion for the program, but aside from the National Science Foundation, it does not include authorizations for the individual agencies in the program.

That's why Senator Johnston, Senator McClure, and I introduced S. 1976. It outlines the Energy Department's part of S. 1067. You see, the Commerce Committee has jurisdiction over the National Science Foundation and NASA, but not over DOE and DOD. Thus, when the Commerce Committee reports S. 1067 it will include authorizations for NASA and NSF, but not for DOE and DOD. The DOE authorizations will be in S. 1976, and I am working with Senator Bingaman and Senator Nunn and others on the Armed Services Committee to include an authorization for a Defense Department high-performance computing program in this year's Defense Department authorization bill.

Each of the committees are taking care of their part to create the whole. I like to think of this as parallel processing. After both S. 1067 and S. 1976 have been reported, we plan to merge the two bills on the Senate floor and send a comprehensive bill on to the House.

Of course, in addition to providing coordination and authorizations for the program, we will also need appropriations to fund it. So, I and my colleagues are working and will continue to work with the Appropriations Committee to get the dollars that will be needed.

Just as the various Federal agencies are going to have to work together, the various Senate committees and subcommittees will have to team up to make this program a

reality. Congressional committees are not known for their ability to cooperate, but more and more often we are facing big scientific problems that require teamwork. Global change research is just one example of a scientific problem that is too big for one agency and one Congressional committee to handle all by itself. The situation is similar with high-performance computing. There are hundreds of talented people in several different agencies, all who have key roles to play in developing and utilizing more powerful supercomputers.

The Department of Energy has a particularly important part to play in a federal high-performance computing program. For years, DOE has been at the forefront in using supercomputers. They had the first Cray and now they have dozens of supercomputers in use helping DOE scientists and engineers design nuclear weapons; predict global warming; build better, more efficient engines; and create advanced materials. DOE scientists are world leaders in developing software to take full advantage of Cray supercomputers, massively parallel supercomputers, and other leading-edge machines. At a hearing on S. 1067 last July, Dr. Carl Heinz Winkler, who until recently was at DOE's Los Alamos Laboratory, demonstrated some of the visualization technology that he and others have developed for supercomputers. The supercomputer graphics he showed were not only stunning to look at, they allow researchers to see patterns in their data that would never be discovered if the data was just listed as a string of numbers.

The keystone of a national high-performance computing program is the National Research and Education Network, the NREN, which will connect over a million researchers around the country at more than a thousand colleges, universities, and federal and corporate research laboratories. This computer network would be able to transmit a billion bits of data per second -- more than 100 times faster than high-speed networks today. This network would allow computer users in every state in the nation to exchange huge data files and to access supercomputers hundreds of miles away. A doctor in New Mexico could send a CAT scan to a colleague in Boston and get a second opinion in seconds. A geologist in Kentucky looking for mineral deposits could log onto the network and get the latest satellite images from a NASA satellite orbitting the earth.

DARPA and the National Science Foundation will be responsible for creating this network, but the Department of Energy has a key role to play as well. DOE engineers have developed some of the fastest computer networks in the world to link their supercomputers at Los Alamos, and are working to create networks capable of transmitting tens or even hundreds of billions of bits of data every second. S. 1976 would authorize funds to continue and expand such efforts.

Most of this work is unclassified so the technology being developed in DOE labs is being put to good use by American firms. Unlike in the past, DOE is working hard at technology transfer. Indeed, there are several firms in eastern Tennessee that owe their existence to computer

technology developed at DOE's Oak Ridge National Lab. In addition, Oak Ridge and the University of Tennessee have teamed up to form the Science Alliance which is now setting up a high-performance computing program in which DOE scientists, university researchers, and private companies can work together synergistically to find new uses for supercomputing.

Almost half of the money authorized in S. 1976 would be for advanced software for supercomputers and other not-so-supercomputers, so that DOE can build on the strong supercomputing applications program that already exists. Much of that money would not go to computer scientists, but to computational scientists -- scientists in other fields who use supercomputing to solve their research problems. Computational scientists come from many disciplines: astrophysics, geophysics, biochemistry, structural engineering, and aeronautics, to name a few. S. 1976 would promote interdisciplinary work by funding research centers where scientists in one discipline, say meteorology, can work with computer scientists and engineers to develop new ways to use this new technology.

I think it is safe to say that I could spend all afternoon giving you reasons to create a national high-performance computing program. The Science Subcommittee has held three hearings so far on high-performance computing, and we have only examined a few of the thousands of exciting uses for this technology. I know several of the witnesses testifying today, and I know that they share with me a sense

of excitement about supercomputing. I also know that they can give you even more reasons than I can to support S. 1976 and to fund a federal high-performance computing program. So I'll just finish up and let them get on with their testimony.

I would be happy to answer any questions you might have.

Senator FORD. Thank you very much, Senator Gore. I do not believe I have any questions, but what we are attempting to do here is to codify the development or the plan that was developed by the Office of Science and Technology Policy. It is my judgment that if we do not move on with it, we will slip further and further behind. And, as you say, it is there and the highways are open, all we have to do is put a vehicle on it. You talk about the child that comes home from school in the afternoon and has access. Well, I have two families in my family, and the grandchildren come home and they have their computer room, and they do a lot of things.

I have been fortunate enough to be able to help them purchase the equipment, the printouts and all of those things. But I know from what I see here in our offices and other companies and other agencies of the Federal Government, that what they have is just a toy and what they are going to be able to have one of these days is a significant piece of equipment where they can do those things, rather than go to a book and labor over trying to find it and put it together. They can call up the Library of Congress or other avenues of connection, and bring that right in and be able to move right on with their education.

I have a hard time adding two and two and getting three or four and Mrs. Ford took over the checkbook, so I know how important it is to have a calculator and things similar to that. You are always ahead of us, Senator Gore, and I appreciate your testimony today, and hopefully, we can move in the direction that both of us want to move here, and maybe we can get Mr. Bromley to come along.

Senator GORE. Thank you, Mr. Chairman, and my thanks again to the full committee.

Senator FORD. As you can understand this afternoon, other committee hearings are going on, involving several committee members, and the Clean Air Bill is up on the floor, and I apologize that you have to look at me and empty seats, but that is the way it is.

Senator GORE. Thank you, Mr. Chairman.

Senator FORD. Mr. Bromley, after all those kind words, would you like to come forward and see if we could take your testimony and ask you a few questions, and we will go from there.

You may proceed. If you want to highlight your statement, the balance of your statement, as was Senator Gore's, it will be placed in the record as if given.

STATEMENT OF DR. D. ALLAN BROMLEY, DIRECTOR, OFFICE OF SCIENCE AND TECHNOLOGY POLICY

Dr. BROMLEY. Thank you, Mr. Chairman. I very much welcome this opportunity to appear before you to talk about S.1976, the Department of Energy High Performance Computing Act of 1989.

There really has been a long working relationship between the executive and legislative branches on this matter of high performance computing. Frankly, Senator Gore has given you a very eloquent statement to which I entirely subscribe. There is no other single thing that I can think of that would have a greater impact on our economic competitiveness, our educational potential and a great many other areas of our national life. So, I would, at the outset, say that I am very much in support of our moving forward.

We have an opportunity in this nation to do something that really will benefit every one of our citizens and I hope that by working together, we can make these things happen.

Frankly, this is the kind of consultative process that I had hoped to find when I first came to Washington. I look forward to working with you and your colleagues and your staff on this.

Senator FORD. Do you think maybe we ought to stop his statement at that point?

Senator GORE. I think we ought to move on to the next witness.

Senator FORD. You will find me doing that. If I get too caught up in something, I just have to break it once in a while.

Dr. BROMLEY. Well, having said that, let me just again, remind you that back last September, as Senator Gore has indicated, my Office sent to you the report entitled The Federal High Performance Computing Program. This was a report that was produced by a committee of the Federal Coordinating Council for Science, Engineering and Technology, the group that has the somewhat unfortunate acronym, FCCSET, that was created by congressional action at the same time as my office. I think, as Senator Gore does, that this was an excellent statement of the problem and of the opportunities that exist here.

The portion of that report that dealt with the Department of Energy was, in a great many respects, very similar to the authorization spelled out in S.1976. And, in the interest of time, Mr. Chairman, I will not dwell on the legislation, but what I do want to do, in my oral presentation to you, is to highlight a few actions that I have taken in my office and in FCCSET that I think are very supportive of what we are discussing here today, and then I will comment very briefly on the legislation.

If I could turn first to the budget process. Early in the Executive Office budget cycle, the Bush Administration decided that a variety of important science and technology areas that are supported by a number of the different R&D agencies should be coordinated to a greater degree than had been the case before, so that we could determine the breadth and scope of the overall federal activity. As a result of that, OMB held interagency cross-cutting hearings on high performance computing, as they did on a number of other areas.

In the late fall, they asked the principals of the FCCSET Committee on Computer Research and Applications to coordinate the various agency submissions. I think this, in itself, is concrete evidence of the high importance that the administration attaches to this high performance computing initiative.

That is a sort of shorthand, if you will, Mr. Chairman, not only for the supercomputer end, but for the networking and for the software development and the development of the personnel that will make all of this actually work.

In the budget you noted that the President gave very high priority to several new federal initiatives that were grouped under the general rubric of developing advanced technologies. One of these was certainly high performance computing.

Now, I am sure that you and your colleagues have noted that the budget request proposed only a five percent increase in fiscal year 1991. I will submit to you, sir, that in very large measure, this is a

reflection of this effort's late entry into the budget process. I, myself, did not arrive here formally until August and some of the agencies did not have an opportunity to focus high level attention on this topic. I believe that the mechanisms that I am setting in place now will insure that this will receive a much higher level of attention in this next budget cycle.

Let me then mention a few of the things that I have in place that I hope will accomplish just that. I became concerned after a few months in office, that the high performance computing initiative was not moving forward with the sort of momentum and speed that I would have liked.

So in January of this year, I convened an informal meeting of agency heads—deputy directors for R&D—of the agencies that support high performance computing. Our objective in that meeting was to initiate a series of discussions of the critical issues involved to get policymakers' views about what needed to be done—how we should proceed. Indeed, one of my goals is to see if it might not be possible to move forward within our existing budgets to get a start on the program that I foresee us moving to in the future.

We have collected informal working documents in that group and we will be meeting again on March 8th. But, this I emphasize, sir, is an informal discussion group. Our goal is simply to better understand the Federal Government's needs and what we can do within the agencies at the present time, and we will continue that activity.

But more important in a more formal sense, perhaps, in our report to you last September, we stated that it was our intention to put together a private sector advisory group for high performance computing. I am happy to be able to tell you now that the assembly of that group is in process. The chairman has agreed to serve. He is a member of the President's Council of Advisors on Science and Technology, PCAST, and we have a very distinguished group. We are not quite ready to announce the membership of that group, but will be in the very near future. This will be a group that will advise the President and OSTP from the private sector, complementing the work that is being done under the Federal Coordinating Council in this whole area.

Turning then to the Federal Coordinating Council, one of my major goals during my tenure of office will be the restructuring and revitalizing of that council.

I believe that one of the problems that has existed in the past, that has made some of its activities perhaps less effective than they might have been, has been the fact that we did not have sufficiently high level participation in the FCCSET itself. I am happy again, to be able to tell you, sir, that in the meetings that we have held so far, and one that I hope to hold in this coming week, we have been extraordinarily fortunate in having participation by a very substantial fraction of the Cabinet. There is enthusiasm and support for moving forward to do the kind of coordinating, integrating, cross-cutting analyses in a number of areas. High performance computing is one, biotechnology is another, material science is another, human resources and training is another. We would like to be able to come to you with a total overview of what the Federal Government is doing in these areas and with a coordinated, inte-

grated plan for the agency participation in the ongoing program in those areas.

I think a good example of what can be accomplished in this area, sir, is the work that has been done during the past 18 months by the FCCSET Committee on Earth Sciences. They have done really outstanding work under the leadership of Dallas Peck from the Department of the Interior in providing the scientific and technical backup for our federal policy making in the area of global change.

I think that in the Computer Research and Applications Committee area, specifically, I would propose to take the following actions. The agenda of the Computer Research and Applications Committee will be broadened to include all aspects of information, science and technology. A Policy Level Committee, composed of assistant secretaries, deputy directors of the appropriate agencies, will be coordinating the fiscal year 1992 budget process for high performance computing, so that we can take forward to OMB and to the President, an integrated program. We were able to do that this year in the area of global change research and, as you will have noted in the budget, we were very successful in that the President has proposed a 57 percent increase in that particular area of the budget.

I also propose to retain the existing high performance computing working group to maintain interagency coordination of these programs. And, let me also say, sir, that I have taken a number of opportunities over this past several months to speak to technical, educational, political groups on this topic, and I have taken the same point of view that I have reported to you here today. I do believe that this is a very important area, one that can have an impact on almost every area of our national life.

We are still very much, as Senator Gore said, in the early stages of the information revolution. We have, as yet, only begun to see the impact that high performance computing can have on our economic competitiveness, on our national security, on education—a particular target because technology and education has fallen far behind most other areas in our society.

I believe that the federal program for high performance computing will very soon, if it has not indeed, already, involved every agency of government and will become a national effort that extends this technology and its benefits as widely as possible.

I want to assure you that my recent actions in OSTP are fully supportive of the goals that I believe we share in this particular area.

But let me turn then, if I may, sir, finally, to the specific legislation in S.1976. As I mentioned before, the information age really is in its infancy. I have said in other contexts that the morning you wake up to find that your toaster is smarter than you are, is going to have a dramatic impact on your view of much of society. The technology upon which information science is based is not just changing faster than any other technology. The rate of that change is accelerating. There is no other area changing as rapidly. We are having trouble even now, predicting how our culture will use this technology in two years, let alone five. And, this problem—it is really not a problem, but a fact—of a very rapidly changing technology, in my view, presents a problem that I would like to bring to your attention with respect to S. 1976.

I believe that any plan that locks-in a program with specific agencies could prevent, could make it difficult for us to have the flexibility that I believe we will need in the out years, and similarly, sir, I believe that specifying budgets for the individual agencies over a five year period in advance could lead again to a rigidity that is perhaps not entirely conducive to furthering the goals that again, I think, we completely share.

For this reason and for these reasons only, I would prefer that we proceed with a broad-based plan that maintains the flexibility and allows the federal agencies to evaluate their priorities within overall coordination on an annual basis. The legislative process we followed in the past has the virtue of providing a regular forum for exchange of ideas between the two branches, legislative and executive. No one piece of proposed legislation for high performance computing is as important, in my view, as getting the job done. We have a shared responsibility for that, and I am persuaded, sir, that by working together, we are on the way and that we are in the process of achieving the goals that we share.

And, with that, sir, I would complete my oral testimony and would be happy to try to respond to your questions.

[The prepared statement of Dr. Bromley follows:]

Testimony
of
D. Allan Bromley
Assistant to the President for Science and Technology, and
Director of the Office of Science and Technology Policy

before the

Subcommittee on Energy, Research and Development
of the
United States Senate Committee on Energy and Natural Resources

2:00 p.m., Tuesday, March 6, 1990
Room SD-366, Dirksen Senate Office Building
Washington D.C.

S. 1976
"Department of Energy High-Performance Computing Act of 1989"

Mr. Chairman and members of the Committee:

Thank you for giving me the opportunity to present the Administration's views on S. 1976, the "Department of Energy High-Performance Computing Act of 1989". There has long been a close working relationship between the executive and the legislative branches of government on high performance computing. I welcome the opportunity to continue and strengthen, that relationship.

As you know, on September 8, 1989, I sent to the Congress a report entitled, The Federal High Performance Computing Program, produced by the Computer Research and Applications committee of the Federal Coordinating Council on Science, Engineering and Technology (FCCSET). The report addressed the implementation plans for most technical departments and agencies of the Federal government for the next five years. The portion of that report which dealt with the Department of Energy was in most respects similar to the authorization spelled out in S. 1976. In the interests of time, I will not dwell at length on this legislation today. But I do want to inform you about five actions that I, and my office, have taken since the issuance of the FCCSET report that I believe are supportive of this general government wide effort. Finally, I will comment on the proposed legislation.

Budget Process

Early in the Executive Office Budget Cycle, this Administration decided that a variety of important science and technology areas supported by a number of different R&D agencies should be co-ordinated to determine the breadth and scope of Federal activities. As a result OMB held interagency cross-cutting hearings on high performance computing – as well as a number of other areas – in the late fall and asked the principals of the FCCSET Committee on Computer Research and Applications to coordinate the various agency submissions. This act, in and of itself, was a first for high performance computing and demonstrates, I believe, that the Executive Office is taking this program seriously. In addition, the President's budget proposal gives a high priority to "Investment in the Future" and, in this category proposes several new Federal initiatives under the heading of "Developing Advanced Technologies". One of those initiatives is High Performance Computing: again, indicative of the importance accorded this topic by the Bush Administration. The budget request proposes a 5 percent increase for FY 1991, in large measure, I believe, because of this effort's late entry into the budget process and because some agencies did not focus high-level attention on this topic.

High-level Meetings

To ensure that High Performance Computing receives adequate high level attention this year, I convened a meeting in January of agency heads and their deputy directors of R&D agencies that support High Performance Computing. Our objective was to initiate a discussion of critical issues, to solicit policy makers' views about what needs to be done, and to determine if there is a consensus among the participants as the basis for extending the Federal program into a National Program. Follow-up information has been collected about the needs and requirements of each agency in the area of high performance computing to meet its mission requirements. The second meeting of the group will be held on March 8. I anticipate that this group will assist the policy and budget coordination process as it develops during the year.

Appointment of an Industry/Academic Advisory Board

In accord with the September 8, 1989 report, OSTP is in the process of appointing an academic-industry panel with a chairman who is a member of the President's Council of Advisors for Science and Technology (PCAST). This panel will provide OSTP and the Committee with an independent assessment of the progress of the program in accomplishing its objectives, continued relevance of the program goals over time, overall balance among the program components, and success in strengthening U.S. leadership in high performance computing, and integration of these technologies into the mainstream of U.S. science and industry.

Reorganization of the FCCSET and the FCCSET committees

The Federal Coordinating Council on Science, Engineering and Technology (FCCSET), established by Congressional action in the enabling legislation for OSTP, is a very effective mechanism for coordinating programs across Federal agencies and for identifying important issues needing broad consideration. The interagency committees and working groups established under FCCSET allow agencies to carry out programmatic initiatives which cross agency boundaries and hence to establish Federal programs which are much broader than would be possible within any one agency or department. The Committee on Earth Sciences R&D program for Global change is an excellent example of what may be accomplished by such a committee.

However, some of the FCCSET committees are not as effective as others. I have therefore initiated a series of actions to restructure and revitalize FCCSET and all of the FCCSET committees. For the Computer Research and Applications Committee, for example, I am taking action as follows:

- The agenda will be broadened to include all aspects of information science and technology.
- A policy-level committee composed of assistant secretaries and deputy directors will be established for coordinating the FY 1992 budget process for High Performance Computing (HPC).
- The existing HPC working group will be maintained for the interagency coordination of programs.

I believe these actions will allow the FCCSET Computer Research and Applications Committee to be far more effective in coordinating budgets and implementing the recommendations of The Federal High Performance Computing Program.

Speaking Out for High Performance Computing

Over the past five months I have had the opportunity to speak to a variety of groups from technical, educational, and political communities. Most of these groups want to know the direction in which OSTP is moving and what our priorities will be. I tell them that the priorities will change as opportunities and problems present themselves, but that, at the moment, my office is focusing on four main themes: Education, Competitiveness, Global Change, and High Performance Computing.

The reason for such emphasis on High Performance Computing is that I believe it plays a critical role in our nation's future. We are still very much in the early days of the information age and have only just begun to develop the broad-ranging applications of computing that will, eventually, transform industry, education, and virtually every sector of our economy, bringing higher productivity and enhanced competitiveness. Thus I believe that the Federal program for High Performance Computing should grow to involve virtually every agency of the government and should become a national effort that extends the benefits of this technology as widely as possible.

I hope that my review of recent OSTP actions assures you of my commitment to supporting the High Performance Computing Program and of my intention to continue OSTP's active coordinating role in the implementation of this activity within the Federal R&D agencies.

S. 1976

As I mentioned before, the information age is in its infancy. The technology upon which information science is based is not just changing faster than almost any other technology, it is accelerating! Change is occurring so fast that we are having trouble predicting how our culture will use the technology in two years, let alone five.

This fast changing technology presents a problem when legislation such as S. 1976 must be considered. Any plan which "locks in" programs within specific agencies could prevent flexibility in out-years. Likewise specifying budgets for individual agencies five years in advance may lead to misallocated resources as circumstances and opportunities change, as they surely will. For this reason, I would prefer that we proceed with a broad-based plan which maintains flexibility and allows Federal agencies to evaluate their priorities for High Performance Computing annually.

The legislative process we have followed in the past has the virtue of providing a regular forum for exchange of ideas between the legislative and the executive branch. No one piece of proposed legislation for high performance computing is as important as getting the job done. I am persuaded that, together, we are in the process of achieving that goal.

Senator FORD. Thank you very much, Mr. Bromley.

We have been joined by my friend, the distinguished senator from New Mexico, Senator Domenici. Pete, do you have an opening statement you would like to make or submit for the record before we get into questions of Mr. Bromley.

Senator DOMENICI. Mr. Chairman, I would like to make a few observations and then I will yield back to you for questions, and then I might have some of my own.

Senator FORD. You may proceed.

Senator DOMENICI. Mr. Chairman, frankly, I am impressed with the notion, the idea that we are trying to develop here. I am hopeful that before we are finished that we will be able to make it a little more specific in terms of what we are trying to do and how we are going to run it.

Let me suggest, that this is not the first time I have had an opportunity to do this. About six years ago, Mr. Chairman, in the State of New Mexico, I worked with some people in New Mexico, and believe it or not, we established a non-profit corporation called Technet that was the first in the country, still is unique. What we attempted to do was to tie the information base that was contained in the national laboratories and the universities to an access system between researchers, private and public, and frankly, we first thought of it in terms of trying to actually own a fiber optics network. And interestingly enough, we did not have to do that, Dr. Bromley.

What happened was we thought that we might put a fiber optics network down the middle of New Mexico, and we did not need it all, you know a fiber optics cable this big carries enough information, we just need a little tiny piece of it. Well, what actually happened is very interesting. The Bell Telephone Company that operates in that part of the world decided that they surely wanted to be the owner of the fiber optic system, so believe it or not, well in advance of when they had planned, they built it and we just leased a little piece of it. So, we have running through about one-third of New Mexico, from the extreme south to the Los Alamos National Laboratory, access to a fiber optics network, and we have now hooked together the universities that are in that corridor and the national laboratories.

The reason I mention this story, this set of facts, is because interestingly enough, I see some similar words here, economic development, technology transfer. Frankly, believe it or not, this foundation is eminently successful and making a whole bunch of money. But, it is merely furnishing a service. For instance, New Mexico has the most advanced method of issuing new license plates for automobiles. You use to have to go wait around. Now if the automobile dealer hooks up to this network, this network for a charge to the state, issues the new licenses for new cars, so we do not wait around forever.

All the lawyers in the state want to access what is going on in the state capitol, like the Corporation Commission. The State of New Mexico said you do it, we sell it. They just sit in their office and punch in the corporate files. They do not have to traipse off to the state capitol to see who the corporate officers are. It is growing.

In addition, the researchers from one university way in the southern half of the state, can talk to a researcher at Los Alamos, just punch in and if they are talking something in genetics, they talk to each other.

But that was pretty easy to manage. There was an entity, a foundation, non-profit run, by more than half business people, the rest were public people. And, I guess, Dr. Bromley, I am all in favor of moving things ahead. I have some confusion in my mind as to how this idea would work. So, before we are finished, I hope I can understand.

Supercomputing is obviously an evolving technology. Again, we will hear from New Mexico on this, because I think it is patent to the world that the demand of Los Alamos National Scientific Laboratory for this insatiable need for rational information was the demand that caused CRAY to come into existence. It was working together with this demand that brought CRAY, and CRAY put their first computers there. Actually there is more computing capacity at Los Alamos Scientific Laboratory than any spot on the globe.

And, I guess I am a little confused and I would like to get it cleared up. Are we trying to promote the next, and the next generation of supercomputers in this legislation or are we trying to make the information that supercomputers yield, not the least of which is an enormous capacity to simulate, as you well know, which is a rather fantastic new art, new science, available? Are we trying to make cutting edge gains or are we trying to share the information capability, the computing capability? I do not quite know.

And, then last but not least, frankly, it is very difficult, Dr. Bromley, to do what we must in all of these ventures, and that is to find a way to bring the private sector in early enough so that you tailor what you are doing to their ultimate needs or any language that says technology transfer or economic growth is very much shortchanged.

Sometimes you wait too long and then it is irrelevant to them. Sometimes you wait too long and it is moving in the wrong direction. So, I would be very interested in how we see all this working together in that capacity and in that regard, and I thank you, Mr. Chairman.

Senator FORD. Thank you Senator. No one has worked any harder and understands this any better than you do, and I am very pleased that you are here today, so that New Mexico will have its input.

Mr. Bromley, let me ask you a few questions about the legislation, so that we can begin to narrow, not what we are trying to get to, but how we are going to put the framework together.

How many of the activities authorized by this legislation are already taking place?

Dr. BROMLEY. It is difficult, Mr. Chairman, for me to give you a precise number on that, but I think that we are touching on all of the points covered by the legislation. Let me, if I may, respond in part to your question, and perhaps in a little more detail to Senator Domenici's remarks, and remind you of what we have underway at the moment.

There are four key areas, as I see it, if we are to move forward, and to be specific to answer Senator Domenici, I believe that we are trying to do all four of these things simultaneously.

The first is that we are trying to make available to a much broader cross section of the nation that now has access to high performance computing the sorts of things that that makes possible. Particularly, this is true of small industries. It is true of schools in particular. It is true of researchers and universities across the country, and I think that the key thing is that by making that access available, the market will expand so rapidly that we will be able to make an adequate base for a viable supercomputer industry in the United States. That is very important.

But, we could, perhaps, go back later to some of the really remarkable, new developments that are made possible by supercomputing. But, let me emphasize the extent to which I agree, both with Senator Gore and Senator Domenici, that one of the key issues here is making this access available with the kind of speed and capacity that allows the full capability.

The government does not own fibers. At the moment, we work with MCI, with Sprint, with AT&T and the other carriers to support the NREN network. This is a national cooperative effort and it is very much like the Technet that the Senator has just described. I would congratulate him on it, because that really was a pioneering venture. We have then the need for the network that makes the capability available. Hopefully that (and I believe it will) will provide the base to continue U.S. development and marketing evolution of supercomputers themselves. But, critically important and often forgotten in this whole area, is the fact that the Senator mentioned, that the original, proud owners of supercomputers were people at the national laboratories, were people in the aircraft industry. These are very sophisticated folk who tended to write their own rather arcane programs. And, if you can say anything about supercomputer software, you will probably have to say that it is not user friendly. If we are going to make this broadly available, there is a real challenge to make software that allows people without extensive background or training to actually use this capability.

And, lastly, Mr. Chairman, we have a serious problem in this country that is highlighted by a single statistic that I have recently come across. That statistic is that in the last four years, interest on the part of freshmen—young American youth entering the nation's universities and colleges—in computer science, among that group of young people, has dropped by two-thirds. That is an important number.

Senator DOMENICI. What was that again?

Dr. BROMLEY. It has dropped by two-thirds, Senator, in the last four years. It has been the most rapid drop in interest in any area of science and technology.

Just to provide the context, interest in physical science has dropped a third in the last 10 years. Interest in mathematics has dropped by a quarter in the last seven. Interest in computer science by two-thirds in the last four.

And so, we are faced, I would submit, Mr. Chairman, with a four-part problem challenge. On the one hand, the hardware, and we

still lead the world in that area. We have the possibility of greatly expanding our fiber network. There is, at this moment, a lot of fiber in the ground in this country that is not being utilized. We also must get software in place so that all these people that Senator Gore and I have been talking about, who can use these facilities to such great advantage, can actually do it themselves without having to interface through a whole group of very highly expert software people. Finally, if we are to maintain our lead in these areas, Mr. Chairman, we have to focus on attracting a much larger fraction of our young people into these areas and training them adequately for the challenges ahead.

Senator FORD. Mr. Bromley, I know you got carried away by Senator Domenici's statement, but let me draw you back to my question if I can.

How many of the activities that are authorized by this legislation are already taking place? I mean, what are you doing, what are we doing? We are authorizing certain things and we do not want to be superimposing legislation on some things. I just want to bring it down so that we will know what is going on now, and if we need this or if we do not. I think we are getting into a question that we may get an answer to before the day is over.

Dr. BROMLEY. I cannot give you an exact percentage, sir, but we are starting actively on all of those areas. We have activity underway. We have planned activity in all of this, and so, frankly, I believe that we already have all of the activities covered in the legislation underway.

I think we have a lot further to go than we are at the moment, but we certainly have started in all four areas. We have the mechanisms in place that will not only coordinate what is happening in government, but also very important, will bring in the private sector input though this panel that I mentioned that will be established under PCAST, because it is very important that we not move forward without knowing what the private sector really wants.

Senator FORD. In your opinion then, should the Federal Government simply encourage the private sector to develop a national research network or take a more substantial role?

Dr. BROMLEY. I believe, sir, that this is something that we should continue to discuss in hearings of this kind and in other ways. Right at the moment, I believe that the system is moving forward in an entirely appropriate fashion toward the goals we have talked about. So, at this moment, I believe that the flexibility that exists in handling this as we are now doing it, is important and I would hate to see any of that flexibility lost because of the very rapid rate of change in the technology.

Senator FORD. Well, let me just ask you this, do you think the legislation is needed if we are doing all that?

Dr. BROMLEY. It is the administration's point of view, sir, that——

Senator FORD. Ah-ah.

Dr. BROMLEY [continuing]. That at this moment, the legislation is not required because we are, in fact, doing the things that the legislation intends to accomplish.

Senator FORD. So, you would not venture your opinion, would you?

Dr. BROMLEY. I believe that we are doing, at this time——

Senator FORD. So, you agree with the administration?

Dr. BROMLEY. I do indeed.

Senator FORD. You said we were the leader a few moments ago in your statement. The United States developed the supercomputer. Some say that the Japanese are but a year or two behind us. How extensive is our lead in high performance computing?

Dr. BROMLEY. Our lead is not extensive, sir, but I believe that the people that we have in our high performance computing companies and our supercomputer companies, still have a very significant edge in terms of innovation and creativity. In the networking area, however, I believe that the Japanese, for example, have stolen a significant march on us, because I am told that on an experimental basis, that the Japanese in some of their major cities, are now not only bringing fiber optics into individual homes, but into individual rooms in individual homes, to try to examine just how broadly computing power can be distributed with high performance. It seems to me that that provides concrete evidence that we have little time to waste in moving forward with our activity in this networking area.

In terms of the computers themselves, we still have an edge. In terms of the networking, we do not have an edge. In terms of software, we still have an edge. In terms of people, we probably do not have an edge, certainly in terms of numbers.

Senator FORD. Well, we develop it but we do not put it in the marketplace. That is kind of our problem, like the VCR. You have heard that forever and we are about to wear that out. But, does this legislation, would it accelerate what we are doing?

Dr. BROMLEY. I do not believe so, sir. I think that at the moment the system that we have in place will move this forward in expeditious fashion. The thing that I worry about, as I noted before, is simply that I am afraid of losing the flexibility that I see that might result from the legislation.

Senator FORD. Well, were you or the administration moved along to put the system, as you call it, into place? Would it be fair to say that you moved along a little faster to develop this system because of the legislation that was coming?

Dr. BROMLEY. I do not think that that is correct, sir. I think that in my own case, at least, and let me speak for myself in this and for my associates in OSTP, we are doing everything we can to move this forward, because as I said at the outset, I cannot think of a single other thing that could have a broader impact and perhaps a more effective impact on our competitiveness, our security, on so many areas of major importance to this nation. And, I particularly emphasize education because I think that is somewhere where, if we are going to meet the President's goals that the Governors agreed to in this past week, we are going to have to move to this kind of technology, broadly in education. That is a whole new area of application.

Senator FORD. Well, I met this morning with four very bright, young people here on one of those week trips with Presidential Classroom. In discussing with these bright, young people, I got a little bit of a feeling that they need more than we are giving them,

whether it is at the local level, state level or federal level. If we are going to give them the opportunity that you had and I had, we are going to have to move a little bit faster than we are.

So, I am very determined that we try to move, not overrun the rabbit, as we would say down in West Kentucky, but that we would keep that rabbit running.

You say, I guess, or the plan released calls for an increase spending of $2 billion over the next five years, and you alluded to the budget that was presented this time, there was only an increase of five percent, but we would not have had that if you had not had your sharp-toed shoes on and got in there late. So, this is five years we are talking about, even in your recommendation, yet, no date has been announced by the administration as to when it will begin.

When will the administration begin funding levels required by the plan?

Dr. BROMLEY. My statement there, sir, is that during this coming year, as we work to produce the fiscal year 1992 budget, I believe that the mechanisms that we have put in place in OSTP, in FCCSET and in PCAST, working with the agencies and working with you and with your staff here, can put together a program that will be sufficiently compelling that in the budget that the President submits roughly a year from now, we will have the kind of proposed funding that will enable us to move forward as fast as we reasonably can.

Senator FORD. We are trying to develop, in my own town, a national communication and educational system. Do you think that has promise?

Dr. BROMLEY. I do, indeed.

Senator FORD. There is a favorite verse in the, that I heard the other day, in the Old Testament about staff. It said that Moses leaned on his staff and died. So, I am finding lots of problems as to where we are going. I want to try to be up here rather than leaning on the staff.

Senator DOMENICI. Were you referring to me?

Senator FORD. No, I was referring to myself. You are way ahead of me. I am just a country boy trying to get along and learning some things here.

Dr. BROMLEY. May I emphasize one point, Mr. Chairman? The Executive Office plan is a five year signal, if you will, to the industry. It is not intended as a funding commitment by the administration.

Senator FORD. Well, you ought to make that very clear. You know, you have to spell it out every once in awhile for us to read it correctly. When it is simple, we have trouble with it.

Dr. BROMLEY. On January 29th, Mr. Chairman, Dr. Huray and I wrote an article for a special addition of the Roll Call Newspaper which dealt with communications, and if I can just quote one sentence from that article. We said, "It is important to send the proper signals to industry about where the U.S. Government is headed and these signals allow industry to make long term decisions about their investments and provide constructive resonance between their plans and government actions." And, I would assert, sir, that this is an action that falls in the area of protrust rather than antitrust, it is not industrial policy, and it is something that

this administration supports strongly. I believe that we are moving forward aggressively on this. I will do everything I can as will my office, to make sure that this program moves forward effectively and successfully.

Senator FORD. Let me tell you two things about what you have just said. One, you do not know when the administration is going to change its mind, that is one, and shift it to something else, so you cannot guarantee it. So, it must be a law in my opinion.

And, you have in your mind that you are going to get the funding for this and that ain't necessarily so. So, on those two points, in my opinion, business cannot depend on what you say. That is just your forecast. That is your desire. That is what you want to do. That is one of the reasons we find ourselves behind Japan because they work very closely with government, private enterprise and employees and it is coordinated. We make a statement today and you read the polls tonight and hell, you change your mind in the morning.

So we just cannot, business cannot depend on government unless it is in the law. That is one reason I think that it might be better for us to begin to focus, find something we can live with and move forward with an authorization, so that business, and industry and universities, and so forth can understand where we are coming from and what they can expect. I am not convinced that you are going to be locked in under those circumstances, because you can be as flexible as you want to, you can change your mind every day, and I do not think you will, but I am telling you in government, we are suspect. Now, that you have joined us, you have become suspect.

I want you to know that, Mr. Bromley, the way people look at you has changed considerably. So, I yield to my good friend here to ask what questions he might have.

Senator DOMENICI. I know we have a lot of witnesses. Both the Chairman and I obviously have many conflicting demands on our time here, but Dr. Bromley, let me ask you, have you seen the legislation, have you reviewed the legislation that is the subject matter of the hearing?

Dr. BROMLEY. Yes.

Senator DOMENICI. Do you support it as a piece of legislation? Do you want the Congress to pass it?

Dr. BROMLEY. My feeling, sir, is that I totally support the goals of the legislation and my only worry, as I indicated, is that I am worried that because this is a field where the technology is changing so rapidly, that if we have legislation that locks us in for this five year period, that we may lose some advantages that would otherwise be ours. It is as simple as that.

Senator DOMENICI. Dr. Bromley, if you have read the bills, let me ask you—the legislation attempts to establish a high performance computing network. Is that a fair statement? I have been trying to put it in a few words. Is that fair as to one of its purposes?

Dr. BROMLEY. It is one of the goals.

Senator DOMENICI. If we are attempting to do that, it seems to me a high performance computing network requires a source of information and a means of transmitting it, and if you have sources and some needs hooked to a transmission system, you have a net-

work. It may not be the network we want, but that is sort of a network. So, let me ask, if that is the case, is the problem in the United States today the actual capability to tie these sources of computing together? If that is the case, we are talking about fiber optics or the like, which is a physical thing. You tie one part of the country to the another with something like that and you permit its access by sources. Do we have a network in existence capable of doing that, and if we do not, do we intend to build one here or to borrow one or to rent one, or how do you see it?

Dr. BROMLEY. We have, if you like, I think one of the best analogies, Senator, in the one that suggests that our present networks are what the super highway system would be like if we remove most of the exit and entrance ramps. We have the major links that tie major facilities, such as the national laboratories, large industries, major research universities—those are tied in—and for those institutions, things work remarkably well. But, what we are missing is the tie, the exit and access ramps that take you to small industries, to local schools, to all sorts of other areas that we have not begun to access yet. That is what we need.

Senator DOMENICI. I clearly come from a very small state, and our experience may be minimal by way of being relevant. But frankly, the goals were precisely what you have just described to make it more than just an interstate. Interestingly enough, the most difficult part was to find nonmajor institutional users that were interested in what was to be accessed. I only suggest that that may be a national problem, and frankly, if you want my opinion, it will be.

Dr. BROMLEY. It will be.

Senator DOMENICI. So, then it seems to me that just on this line of questions, do you know whether or not the sources of this fantastic capability that already exists, and we are trying to cause more to come into being, do you know whether or not legally, these sources, aside from national security information, are currently in a legal status of availability? What if you have this network?

Dr. BROMLEY. Certainly, a very large fraction of the material, Senator, is available, publicly available to any user who wants to tie in. There obviously are some aspects that, for various reasons of classification, because of national security, because of proprietary rights, are not available widely, but that I really believe, sir, is a very small fraction.

Senator DOMENICI. So, at least we ought to inquire of those that have some of this institutional capacity that is major, such as Los Alamos, Livermore, Oak Ridge, MIT and the like, University of Pittsburgh, whether their information would be accessible and under what conditions?

Dr. BROMLEY. Yes. And I believe that the answer is at least partially in, that many of these institutions have been very enthusiastic about stating that they are more than willing to make their information accessible and that they are more than willing to assist in making it available in a user friendly fashion to people who will use it. You are absolutely right, sir, that one of the problems is convincing people that they should not be afraid of this. That, in fact, this whole new capability is something that will really be

useful to them. There is a natural reluctance to tackle something that is so totally new.

Senator DOMENICI. Well, my last question and clearly Mr. Chairman, we have got to get on with the other witnesses, but it does appear to me a subject only to whether the experts who follow you that have been dealing with the supercomputers and doing incredible things with them, agree. But I think one great challenge that we ought to try to meet is to make the computing, supercomputing or high performance computing, to borrow your words, user friendly. I think that is a challenge. I am quite sure that they are going to testify that there is not as broad a use for the product today as they might envision, and that probably requires moving some software capability around and forward.

My last question to you, Doctor, is I hear you talking about you getting some of these things done. Is one of the jobs that Dr. Gene Wong, nominee for associate director, is he going to be kind of your lead person in this regard?

Dr. BROMLEY. Very much so, Senator. Professor Wong had his confirmation hearings last week and I very much hope that he will be confirmed. He comes, as you know, from many years as Chairman of the combined Departments of Electrical Engineering and Computer Science at Berkley. He will be, if confirmed, my Associate Director for Physical Science and Engineering and he has, of course, particular expertise in this field and has assumed major responsibility for following up and making this happen. But, let me also, sir, say that another of my associate directors, who has been, again, through his confirmation hearing last week, is Dr. William Phillips, who has an extensive industrial background in Dupont and Mellinkrot as a senior research and development executive, and he has been working most recently with the Governor of Missouri setting up computer-linked, high technology centers. They will be working together with me and other people on my staff and I think we do have the background to move this forward.

Senator DOMENICI. Thank you, Mr. Chairman. Thank you Dr. Bromley.

Senator FORD. Thank you. We have been joined now by another senator from New Mexico, Senator Bingaman. Senator, do you wish to make a statement.

Senator BINGAMAN. Mr. Chairman, I will just put a statement in the record and compliment Dr. Bromley. I know of his views on this important subject, and I will not hold up your hearing at this point. I look forward to the additional testimony.

Senator FORD. Thank you very much, Senator, and your statement will be included in the record as if given.

Dr. Bromley, we thank you and look forward to working with you in the future.

The next witness this afternoon will be Dr. David B. Nelson, Executive Director, Office of Energy Research. Welcome, Dr. Nelson, and you may highlight your statement. It will be included in the record and you may proceed.

STATEMENT OF DR. DAVID B. NELSON, EXECUTIVE DIRECTOR, OFFICE OF ENERGY RESEARCH, DEPARTMENT OF ENERGY

Dr. NELSON. Thank you, Mr. Chairman, Senators. I appreciate the opportunity to represent the Department of Energy in today's hearings on S. 1976.

By way of background, as Executive Director of the Office of Energy Research, I have responsibility for the Department's basic research programs. I am also the Director of Scientific Computing in the Department of Energy. I am responsible for most of the activities explicitly described in S. 1976. I am also the chairman of the Committee on Computer Research and Applications of FCCSET that Dr. Bromley just described, and finally, I led the effort under this committee that drafted the OSTP report, Federal High Performance Computing Program.

I would say that this report is a blueprint for a coordinated, multi-agency program in high performance computing, and that its role for DOE is closely parallel to many of the provisions of S. 1976.

Now, in my oral comments, I will merely address three questions and leave the rest of my testimony to the written record.

These questions are, one, why is the Department of Energy interested in high performance computing. Two, what is the Department's role in the Federal High Performance Computing Program, and three, what are the Department's views on S. 1976.

Let's take these one at a time. First question, why is the Department interested in high performance computing? The answer is simply that many engineering, technical and scientific problems that are critical to our missions require the most powerful computing systems that we can get, and in many cases, what is available today is still not nearly sufficient for the magnitude of the problems that we need to solve.

Examples of these problems under the Department's mission include: the development of energy sources, such as nuclear fusion and fission, solar energy and fossil fuels, the development of new materials, such as superconductors, ceramics and metal alloys, understanding of biological processes, such as protein structure and the human genetic code, development of better and more reliable defense systems, understanding global climate change, and finally, probing the fundamental structure of matter in our high energy and nuclear physics programs.

Since the beginning of the electronic computing age, over 40 years ago, DOE and its predecessor agencies have been at the leading edge of computing. We learned early on to use computing as an important tool in parallel with experiment and with theory, to solve scientific problems of the sort I described above.

DOE is by far the largest user of supercomputers in the world. However, the way in which we use computers has changed over the years. In the 1950s and 1960s, we installed supercomputers in a few of our national laboratories and made them available, primarily for classified calculations, to the researchers locally at those laboratories. In the 1970s, we realized the importance of making supercomputers widely available to our researchers in laboratories and universities and we developed a nationwide supercomputer data network that linked users to remote supercomputers. This network is

one of the forerunners, one of the pilots, that gives us confidence that we can tackle the NREN, the National Research and Education Network.

In the 1980s, we added work stations, new types of computer architectures, and software to link all of these together into a more productive environment. This environment we call high performance computing.

Now, for the 1990s, we are working with other agencies and with the computer and communications industries to develop the much more powerful computing environments that we will need.

In the next few years, we can confidently expect to achieve 1000 times greater computing power than is available today. We are cooperating with other agencies in high performance computing to harness this power for our own missions, and we are working with American industry to bring the benefits of high performance computing to the economy generally.

Let me turn, now, to the second question. What is the Department's role in the Federal High Performance Computing Program? The short answer is that the Department's ongoing research programs already provide a template upon which we can grow in the roles that are assigned to the Department of Energy Federal High Performance Computing Program.

Mr. Chairman, you asked the question, what are we doing today in the Department of Energy that is relevant to S. 1976? If I could use an analogy, we have runners on all the bases, but some of those runners run a little slowly. That is, we have the basis for all of the things that are called for in the Federal High Performance Computing Program, and therefore, we have some confidence that we can carry out the activities that are called for in that Program, but we certainly do not have, at this point, the full capability that is envisaged in that Program.

Dr. Bromley described the four parts of the National High Performance Computing Program, namely, the National Research and Education Network high performance computer system development, and computer software and application development, and basic research human resources.

In my written testimony, I give a number of examples of the Department's activities in each of these four areas. However, for my oral comments, let me just touch briefly on technology transfer, because it is an item that is covered specifically in S. 1976 and is an important part of the Federal High Performance Computing Program. I would say, though, that the technology transfer of the Federal Government and the Department of Energy is an evolving activity. It is not yet one that we have mastered.

I think that the Federal High Performance Computing Program and the Department's role clearly envisages a close cooperation with industry in order to bring the benefits of high performance computing out of the laboratories, out of academia and into the economic marketplace at a far more rapid pace than has been the case in the past. I do not think we can afford to do otherwise.

Let me turn to my third question. What are the Department's views on S. 1976? First let me say that the Department is eager to cooperate with Congress in developing and utilizing high perform-

ance computing. I think you understand that this is an important tool for improving productivity in science and industry.

However, as Dr. Bromley said, it is also a rapidly evolving technology that may be ill-suited to legislative permanence. In particular, we do not believe that legislation at this time is necessary to carry out the Federal High Performance Computing Program.

In fact, it might impede flexibility of the agency's response and coordination of this important topic. I can assure you that Secretary Watkins has taken a personal interest in the Department's activities in high performance computing. And he intends that the Department maintain its lead role in this area.

As the Senate contemplates high performance computing legislation, I ask that you be mindful of the Federal High Performance Computing Program, and the fact that this Program builds on the strengths of several agencies, including DOE, DOD, NASA, NSF, Commerce and HHS, in order to achieve a comprehensive program that would be beyond the attainment of any one agency. In my opinion, any legislation that would limit or distort this comprehensiveness would be counterproductive.

This concludes my remarks. I would be happy to answer any questions that you might have.

[The prepared statement of Dr. Nelson follows:]

Statement of David B. Nelson

Executive Director

Office of Energy Research

Department of Energy

before the

Subcommittee on Energy Research and Development

of the

Committee on Energy and Natural Resources

United States Senate

March 6, 1990

Mr. Chairman and Members of the Subcommittee:

I am pleased to be here today to discuss S. 1976 entitled, "Department of Energy High-Performance Computing Act of 1989." This bill aims to provide continued United States leadership in high performance computing. As Director of the Scientific Computing Staff for the Office of Energy Research, I am responsible for most of the activities covered by this bill. I am also the chairman of the Committee on Computer Research and Applications of the Federal Coordinating Council of Science, Engineering and Technology (FCCSET). I also led the effort under this committee that drafted the Office of Science and Technology Policy (OSTP) report entitled, "The Federal High Performance Computing Program." This report, released on September 8, 1989, contains the basic ingredients for the High Performance Computing Plan called for in Section 3 of S. 1976.

S. 1976 authorizes the Secretary of Energy to: (1) develop a long-term strategy for research, development, and application of high performance computing; (2) implement that strategy in conjunction with other Federal agencies as part of a "National High Performance Computing Program;" and (3) ensure the appropriate transfer of high performance computing technology to United States industry. I am pleased to say that S. 1976, the OSTP report, and the Department's ongoing research programs in high performance computing are all aimed at the same goals. All recognize the need for continued progress in this vital area, progress that will have important impacts in broad industrial sectors, including electronics, telecommunications, aerospace, energy, transportation, materials and chemicals, and health.

From the early days of scientific computing, the Department of Energy and its predecessor agencies (the Atomic Energy Commission (AEC) and the Energy Research and Development Administration (ERDA)) have played a key role in the development and application of supercomputers to scientific computation and modeling.

Originally, the AEC used supercomputers to further its mission in the design and development of nuclear weapons, a challenging set of complex problems that do not lend themselves to simple analytical methods or to detailed experimental measurements that can be repeated often and economically. The first digital computer that resembles those familiar today, known as the ENIAC, was built at the University of Pennsylvania under a contract with the U. S. Army's Ballistic Research Laboratory in 1943 and used by Los Alamos scientists to model the implosion concept for triggering a nuclear device. Subsequent collaborations with supercomputer vendors, such as IBM, Univac (now Unisys), Control Data Corporation, and Cray Research produced an industry and a cadre of scientists capable of leading the world in computational science applications.

The establishment in the Department of Energy of the National Magnetic Fusion Energy Computer Center at Lawrence Livermore National Laboratory in 1976 originated in the need to supply supercomputer capability to geographically dispersed researchers in the magnetic fusion energy program. Today, that center has expanded to serve all programs supported by the Office of Energy Research - some 4,000 - at universities and laboratories nationwide. This

service includes a modern data network to provide the highest quality access to state-of-the-art supercomputers and the appropriate scientific software.

I would like to comment briefly on DOE activities related to the Federal High Performance Computing Program. The plan lays out the need for a National Research and Education Network (NREN). The Department is an active member of the newly established Federal Network Council (FNC) and other committees associated with the NREN and is committed to working with other agencies to meet common research network needs to the extent that mission needs allow.

The report describes the role that the Department would play in research and development in the field of software technology to support "Grand Challenge" problems. The examples of "Grand Challenges" cited are appropriate for the Department. Several such projects are already underway, including projects in fundamental properties of matter, prediction of global change, modeling of high critical temperature superconductivity, and modeling of protein structures.

The Department is supporting these projects with supercomputer resources at the National Magnetic Fusion Energy Computer Center. In addition, the Department has established experimental high performance computer research groups that have acquired computer systems with architectures that have the potential of becoming efficient systems for future investigation of some of the "Grand Challenges." The DOE national laboratories participating in this program include Ames Laboratory, Argonne National Laboratory, Los Alamos National Laboratory, Lawrence Berkeley Laboratory, Lawrence Livermore

National Laboratory, Oak Ridge National Laboratory, and Sandia National Laboratory at Albuquerque. Also, several major universities, including Illinois, Cal Tech, and New York University, are participating in this program. Most of these research computers are available by remote access through networks.

Typically, these research high performance computers have been acquired at an early stage of commercial development in order to evaluate their potential. The Department often undertakes cooperative activities with the computer manufacturers in order to accelerate the availability of high performance computers for the Department's missions. These activities can include the joint development of system or applications software, the implementation of a range of applications that may assist in attracting a broader market (e. g., software packages for modeling combustion processes or protein structures), and analyzing various technical aspects of system characteristics that may need attention.

The report also describes the roles of the various agencies in supporting research and development of high performance computer systems. The Department has been a leader in acquiring early market and production model computer systems for its research programs, including the first Cray 1 and the first Cray 2. In addition, the Department has been able to acquire early models of experimental parallel computers and to provide support for research in algorithms, languages, and systems for these "future supercomputers," often in combination with other Federal agencies funds or with funds from other sources. For example, Oak Ridge National Laboratory (ORNL) recently acquired

the first of the new Touchstone parallel computers being developed by Intel Scientific Computers, Inc., with partial support from the Defense Advanced Research Projects Agency. The ORNL acquisition was made possible by joint support from my office, other DOE and non-DOE programs at the Laboratory, the University of Tennessee and Vanderbilt University. The Touchstone project has the potential to develop into a true supercomputer system, including up to 2048 processors, each of which is capable of approximately 50 million operations per second.

An important aspect of the report addresses the Basic Research and Education activities, including support of basic research, transfer of technology to industry, and support for educational activities at the national laboratories (including support for postdoctoral fellowships as well as high school teachers and students). The Secretary of Energy is particularly supportive of efforts to improve mathematics and science education at all levels of the educational system. Among our ongoing activities are: (1) the High School Supercomputing Honors Program at the National Magnetic Fusion Energy Computer Center at Lawrence Livermore National Laboratory; (2) the establishment of several distinguished fellowships (named after eminent mathematicians and computational scientists) at the DOE laboratories; and (3) the establishment of the National High School Supercomputer in collaboration with Cray Research that will provide supercomputer time to students in the Honors program after they return to their high schools.

Technology transfer between the Government and the private sector is an important part of the Department's operating philosophy. The limited

availability of proven applications software is a critical bottleneck to the effective use of high performance computing in industry. Many of DOE's research activities generate software as a part of this research, and technology transfer has made much of this software available to the private sector, for example through the National Energy Software Center at Argonne National Laboratory.

As an illustration of the effectiveness of the Department's program in sharing new technology and knowledge with industry, two projects supported by our high performance computing program were selected by "Research and Development Magazine" for two R&D 100 awards in 1989: one for improvements in networking techniques awarded to Lawrence Berkeley Laboratory, and one for "parallel computing software for scientific problems" awarded to Sandia National Laboratory at Albuquerque. To further improve the Department's technology transfer of software we have begun a pilot program that will systematically identify promising software for industrial applications that has been generated by DOE research activities, provide limited funds to improve the reliability and documentation of this software, and encourage its transfer into the private sector.

The Department is eager to coordinate its work on software with other Government agencies by means of the FCCSET Committee on Computer Research and Application. The Department of Energy (including its predecessor agencies) has traditionally been the key agency involved in advanced computing technology and in the use of this technology to fulfill its missions. The Department will continue to be one of the major participants of the high

performance computing program. I believe this is appropriate and desirable if our nation is to maintain its leading position in high performance computing. This is vitally important for our national security, both economic and military, for the development of adequate energy resources, and for the pursuit of new knowledge in the physical and biological sciences.

The OSTP report and the follow-on activities by the FCCSET Committee on Computer Research and Applications are providing a firm foundation for a strong Federal high performance computing program. We do not believe legislation is necessary. In fact, it might restrict the flexibility of agencies' responses and coordination of this important research area.

This concludes my statement. I will be happy to answer any questions you may have.

Senator FORD. Thank you very much, Dr. Nelson.

Does the plan that the Office of Science and Technology Policy released—the plan they released—does that expand or maintain or increase the role of the Department of Energy in this area?

Dr. NELSON. I mentioned that the base for that—for the Department's roles—is already in existence. So the plan does not take the Department of Energy into any new area that it does not currently have experience with. It does expand within those areas the magnitude and the scope of those activities. It is giving those base-runners a bit of training, so they can run those bases a little faster.

And by uniting those components into a comprehensive plan it makes more viable the technology transfer and the work with American industry that I alluded to.

Senator FORD. The Department of Energy receives much credit for having helped Cray become the company it is today. Is the same kind of help given to Cray still available for other companies?

Dr. NELSON. Yes, sir. Let me give you a couple of examples. I think that in a few minutes you will hear from Dr. Hecker, the work that Los Alamos is doing with the Thinking Machines Corporation in helping to develop and test out the Connection Machine. I would mention another example, Senator Domenici and Senator Bingaman, in your State, the work that Sandia Laboratory is doing with the NCUBE Corporation to develop hypercubes. Excuse the jargon, these are another form of very high performance computing.

And finally, Mr. Chairman, in Tennessee close to your State of Kentucky, the work that Oak Ridge National Laboratory is doing with the Intel Corporation in the brand new Touchstone super computer that was developed partly with Defense Advanced Research Projects Agency Funding.

So, yes, sir, the Department of Energy is continuing to work with all promising architectures. One of the things that we value is the broad base of our laboratories and our research interests, so that we can accommodate in our fold, if you will, new architectures as they emerge without having to kick other ones out of the fold.

Senator FORD. I think, Dr. Nelson, you are going to have witnesses that are going to say that Cray is the favorite. And whatever is left over—it is like these letters I am getting from those who oppose—what is it the new tax you all are trying to promote, Pete? Capital gains. It is all crumbs, you know. And so I am getting little packages of crumbs from all over the United States.

Did you get any yet?

Senator DOMENICI. Well, you just develop your own package, and we will give you yours too.

Senator FORD. Of crumbs? [Laughter.]

Senator FORD. But I am saying that we will have that. So I want you to be alert to that, that there will be some.

Dr. NELSON. I thank you, sir.

Senator FORD. You say this legislation is not necessary, but you did not oppose it.

Dr. NELSON. That is correct, sir.

Senator FORD. So, you are neutral?

Dr. NELSON. I am speaking for the Department, when I say the legislation is not necessary. But I do not use, nor does the Depart-

ment, the word "oppose." And I appreciate your noting that distinction.

Senator FORD. All right. That is like the fellow who died, they said he did not look natural, he looked neutral. [Laughter.]

Dr. NELSON. Shades of gray, I guess.

Senator FORD. Yes, shades of gray, yes. I understand the Department will no longer agree to purchase computers before they are built. Is this the case, and if so, do we not need legislation to kind of change that practice?

Dr. NELSON. Well, sir, different people could vary as to the definition of "built." If you have a prototype, or a first of a kind, a lot of companies would not purchase it because they would say, it is not proven technology.

The Department of Energy has, does and will continue to do this. Now, when you get into strictly "paper" machines, then there is a question of judgment. And I would not say that we have an explicit policy, but it does appear that we have our plate pretty full dealing with those concepts that have evolved to the prototype stage.

Senator FORD. You lost me a little bit there somewhere. And I am trying to catch up.

Dr. NELSON. Too many left turns? All I said was that we have——

Senator FORD. Kind of go to the right a little bit, will you. I am conservative.

Dr. NELSON. We have limited resources. And we have to establish priorities. And we believe that our plate is pretty full, purchasing and working with computers that are already at least at the prototype stage. We have no explicit policy against purchasing "paper" machines. But I tell you, I would rather take a prototype today than a "paper" machine tomorrow, unless the performance of that "paper" machine is so much better than the prototype today that I can afford to delay.

Senator FORD. Well, let me ask, how does the Department proposed level of spending for 1991 compare to the funding level the OSTP plan would require?

Dr. NELSON. That is very difficult to answer. The OSTP plan would call for an increase in the Department's funding in the first year of the program in the neighborhood of $30 million.

Given overall Department priorities, the Department does not have the $30 million of additional funding required in this area.

Senator FORD. So you spend $30 million annually?

Dr. NELSON. We spend in the neighborhood of $60 million in direct budget lines that we can relate immediately to high performance computing. If we were to take those budget lines that are justified for other purposes, but which have a very hefty high performance computing component, then it is well over $100 million a year.

Senator FORD. Well, my final question. Since the plan the bill requires is fairly broad, and I believe it is updated annually, how does it restrict the flexibility of the Department? And if it does restrict, what kind of legislation do you need to give you the needed flexibility?

Dr. NELSON. Did you describe the plan, sir, or the bill? I am not sure I understood.

Senator FORD. The bill requires—the plan the bill requires is fairly broad and is updated annually. And I wondered how if it updated annually, how it would restrict your flexibility?

Dr. NELSON. The bill calls for specific activities in specific areas and lays out specific levels of funding in each of those several areas.

As Dr. Bromley has testified, when we look to the future, we cannot be sure that exactly those activities in exactly those areas, with exactly those funding levels are appropriate ones. And legislation once passed tends to hang around. And it is for those reasons that we believe that the legislation is unnecessary and might impede our flexibility to respond to this very rapidly changing technology.

Senator FORD. I will send you some questions in writing on this particular case, and ask you some more points rather than keep us further. So, I will turn it over now to Senator Dominici.

Senator DOMENICI. Mr. Chairman, I wonder if I might ask you with reference to the schedule. We are at Dr. Nelson and then we had a panel number 1 and a panel number 2. I am sure that I want to be here, and I think Senator Bingaman wants to be here when Dr. Hecker testifies.

I am wondering if you could accommodate me and perhaps us by at least calling the second panel along with the first or all as one panel?

Senator FORD. We will put them all at the table, because I have a Kentucky friend here also. And I would not want you to get anything Kentucky did not get. [Laughter.]

Senator FORD. We will just put them at the table at the same time. Will that be all right?

Senator DOMENICI. That is great.

Dr. Nelson, let me go back a minute with you. You described the four goals. And would you repeat them again for me?

Dr. NELSON. Yes, sir. The National Research and Education Network.

Senator DOMENICI. All right. And you say that to the extent that DOE is currently involved in that effort, which anybody that wants to find out what it is, we could find out what it is. You are involved in doing that right now?

Dr. NELSON. Yes, sir. We are a member of the Federal Network Council, and we are already linking our computer network traffic with that of other agencies.

Senator DOMENICI. Right. What was the second point?

Dr. NELSON. High performance computing systems, hardware and system software. That is the second point, sir.

Senator DOMENICI. I thought the word was development.

Dr. NELSON. That could be related, yes. Development of high performance computing systems.

Senator DOMENICI. Okay, so the development of them. Is the Department of Energy doing that?

Dr. NELSON. Yes, sir. The mode that we are doing it is working with vendors in a cooperative fashion. I think maybe Dr. Hecker will describe that a little bit more.

Senator DOMENICI. And then you had software development.

Dr. NELSON. Correct. Software development in particular looking toward what we call the grand challenges, some of those application areas like global climate change and energy sources.

Senator DOMENICI. And then you had education.

Dr. NELSON. And basic research, yes, sir.

Senator DOMENICI. Okay. Might I ask now, one of the strengths we are trying to build as a nation now—if I got this right—is to take other major American Government activities of a science and research type, such as NASA and others, and to maximize the capacity to do these four things by melding them together, hopefully, in a synergistic way so that the nation gets all of these with as many of the super agencies, like NASA and departments like Energy and others, all tied in. Is that correct?

Dr. NELSON. Yes, sir.

Senator DOMENICI. Now it seems to me that one of the difficulties—and maybe it is your flexibility concern—as I gather, you are suggesting to this committee that the national goal of putting them all into a network, for lack of a better word, to do all of these things. That you want to urge that we not make it less effective by limiting the flexibility of its component parts, in particular, in your case, the Department of Energy as we meld them together.

Dr. NELSON. That is correct.

Senator DOMENICI. And I think that is a very serious concern. We have to worry about that.

Now if we are going to put them together in a network, somebody has to, in a sense, assume some overseeing or overseeing role—I am not familiar with the bill, so forgive my ignorance at this point—what does the bill say about that?

Dr. NELSON. The bill, S. 1976, is to a considerable extent silent on who takes the lead role. It is a bill aimed specifically at the Department of Energy. And I think maybe it is appropriate that the drafters did not try to call out the roles of all the other agencies.

I would say, Senator, that the bill as well as the federal High Performance Computing Program envisions a great deal of peer cooperation. That is, a mode that this country maybe has not had so much practice in. But it is one that seems to work pretty well in other places: That you try to form a consensus, a solution that satisfies all the parties needs, and get a little bit away from the dictator mode.

You have somebody who keeps the time, somebody who calls the meetings, and in some cases, somebody who holds the budget. But try to carry out some of these activities in a cooperative mode. We cannot tell industry what to do. And we also think it does not make sense for one agency to tell another agency what to do. We are trying this cooperative mode.

Senator DOMENICI. But it is true that the mission of the parts, singularly DOE, NASA, has a positive impact in that the fact that they are driven has caused some of this to advance based on the momentum of the mission of one of the parts.

Dr. NELSON. Yes, sir.

Senator DOMENICI. So, if we are going to have a referee, or a peer, whatever you call it, you are suggesting—at least tacitly here—that we make sure that the component parts remain driven by their mission, and that as they accomplish it, they contribute.

Dr. NELSON. Absolutely. We see—and I think historically this has been true—that the research and development activities of the federal agencies lead those of the broader economic activities of this country by five to 10 years in this area. And so by focusing on our missions without trying to go toward a welfare or industrial policy mode, we are driving ourselves, but also contributing to that broader economic activity.

The one ingredient we have to watch out for, and what the program calls for, is careful attention to technology transfer. When we are leading, we have somebody that is working with us.

Senator DOMENICI. Okay, you mentioned—this is my last question—you mentioned somewhere in your testimony, not oral but written, that the availability of software is a critical bottleneck to quote, "a critical bottleneck to" industry.

So without knowing about your testimony, or at least reading the question my staff had prepared here, I had said to Dr. Bromley that I could see one of the problems with user-friendliness, the evolution of software.

You are suggesting that is a problem. Does this legislation help solve that problem? Or is there something you would recommend that would?

Dr. NELSON. I think that the agencies, in particular DOE, already have the flexibility to solve that problem. The various technology transfer acts, for example, Stevenson-Wydler and Bayh-Dole, give us the authority—

Senator DOMENICI. Wait a minute. You better mention ours, or you are really in trouble. We just passed one, so now you can add ours.

Dr. NELSON. All right. I certainly know, Senator, of your work in the technology transfer area and cooperation between the laboratories and industry. At any rate, we believe that we have plenty of authority. In fact, as my written testimony says, I have started a pilot project in the Department of Energy to identify software that is of interest to specific industry and that has been developed in our national laboratories or university programs and to put a small amount of federal money into making that software robust so that it works well. And then, working with the private sector, either with computer companies, with software vendors or with industrial users of the software themselves to turn that software into industrial strength user-friendly software.

This is a pilot project. I do not know if it will work. But it is an attempt to deal with that bottleneck. And it is taking place under current legislation.

Senator DOMENICI. Doctor, with reference to computing and the demands of the laboratory driving the technology out, there has been a suggestion that because Los Alamos had such an enormous demand, they have pushed Cray technology to the limits and brought it into being. There has been some suggestion that maybe there is a lack of spreading this out to other prototypes or notions.

Parallel Computing is not the Cray Computing System. Right?

Dr. NELSON. I would not want to speak for John Rollwagen, but let us just say that one type of parallel computing, massively parallel with thousands of parallel processors is not what Cray Research, currently delivers.

Senator DOMENICI. But it is one of the evolving computing technologies. Right?

Dr. NELSON. All American high performance computing manufacturers are going to parallelism, yes.

Senator DOMENICI. Is the Department of Energy pushing that technology?

Dr. NELSON. Yes, sir, we are. We are trying to ride both horses, because we do not know what level of parallelism, how many processors is appropriate. So we are continuing to buy and use Cray computers. And we did buy and use ETA computers. At the same time, we are buying and using computers manufactured by Thinking Machines, by NCUBE, by Intel and by BBN and by other companies.

Senator DOMENICI. Thank you.

Thank you, Mr. Chairman.

Senator FORD. Senator Bingaman.

Senator BINGAMAN. Thank you, Mr. Chairman.

First let me just say that I support Senator Gore's bill. I am cosponsoring it, and I compliment your coordinating counsel for the work that went into this report. I think it is excellent. I have got to say, I am still somewhat nonplussed by the administration's position on this issue.

As I understand it, they take the view that they do not believe legislation is necessary. And secondly, they do not support the money needed to get this proposal implemented. That strikes me as fairly lukewarm support.

My concern on this is that—and I understand the many reasons why we are trying to do it, as you are suggesting, sort of by consensus, and try to get all of those who really have a clear understanding of the importance of this and the various agencies to do their part. And then try to come to Congress each year and ask us to go around to each of the budget recommendations that the administration sent up and plus it up here by $57 million, and plus it up here by something else so that it can happen. But that is a onerous way for us to have to try to do this thing.

It would be much better, in my view, if the administration would make a commitment to this project and would go ahead and recommend to the Congress a reasonable level of funding, and direct that agencies cooperate to get this done.

Do you have any comments on that?

Dr. NELSON. Yes, sir, I do. I can certainly relate to your concerns. And I assure you, we face them ourselves in putting this program plan together among the Federal agencies. I think we all learned a lot from each other and overcame some of our own biases.

Now Dr. Bromley testified a few minutes ago on his efforts to raise the visibility of this issue in the administration and to get a good head start on the 1992 budget. And I guess I would ask for some indulgence until you see what happens as a result of these efforts.

I know that in the Department of Energy Deputy Secretary Henson Moore attended the first meeting on January 8, and he is expected to attend the second meeting on March 8. I also indicated Secretary Watkins' personal interest in this area. He would be going to the meeting too were it not for scheduling conflicts.

So I think that within the Department of Energy we will see in the 1992 budget process and conceivably even in the 1991 outcome a somewhat greater attention to this issue than the President's submission would indicate. I can say that with somewhat more confidence for 1992 than for 1991.

I do not know that having authorization would in and of itself solve the funding problem, because it still has to go through appropriations.

I appreciate your dilemma. And I can say, sir, that we share the same dilemma, and we are doing what we can, and I think fairly productively. So watch what we do.

Senator BINGAMAN. Well, as I say, I favor what you are trying to do. I just hope that the administration will be more supportive of the funding levels in the budgets that they submit over the next several years. Because it is difficult to do what Senator Gore and others are trying to do this year. And that is, go around to each of the appropriate people on Capitol Hill and try to correct or modify the administration's budget, so that something like this plan could take place.

Let me ask about cooperation you have with DOD. One of the things I have been advised of is the Department of Defense needs to put $57 million additional into what has been proposed by the administration, if this is going to happen, in this year's budget.

What is the level of support that you are getting from the Department of Defense?

Dr. NELSON. The Department of Defense, especially through DARPA, is in its budgets and in its activities supporting the program as much as the Department of Energy. In other words, I cannot complain. None of the agencies, as Dr. Bromley testified, has been able in the 1991 budget to come up with significant amounts of additional funding.

However, base funding in this area is significant. We estimate that. We estimate that about $500 million a year goes into activities that are supportive of or related to the High Performance Computing Program plan. Within those activities there is a lot of cooperation between DARPA, and to some extent also the military services, and the Department of Energy.

For example, the Connection Machine that was installed at Los Alamos was developed partly under DARPA funding. And there has been a lot of discussion between DARPA personnel and DOE personnel as to how best to test out and verify those machines. I mentioned that the Intel machine, the Touchstone down at Oak Ridge was also supported partly by DARPA funding.

The Department of Energy and DARPA are cooperating by pooling funds for our common needs--I do not think there is any problems with appropriations laws here—in testing of advanced networking concepts. I could go on with a long list. But I think it is good cooperation.

Senator BINGAMAN. On a slightly different issue, let me ask how—there are two sides roughly speaking in the Department of Energy. You have got the defense and the nondefense. You are the nondefense, and you are the lead person on this project within DOE.

Dr. NELSON. That is correct, sir.

Senator BINGAMAN. And within this coordinating council as I understand it. Is that right?

Dr. NELSON. That is correct, sir.

Senator BINGAMAN. How well is the defense side of DOE involved in the project that you are heading?

Dr. NELSON. Increasingly. To be very candid, it has taken some time to get the attention within Washington of the defense side of the Department of Energy. But I can assure you that attention has been gotten. I can also assure you that the laboratories have helped by talking to people in Washington Defense Programs about the value of high performance computing in the Defense Programs.

So I think the cooperation now is excellent. And we are looking toward more joint programs as we develop the 1992 budget. It could not be better.

Senator BINGAMAN. Thank you.

Senator BINGAMAN. Thank you. That is all I have.

Senator FORD. Thank you, Senator Bingaman.

Dr. Nelson, I always have to look for something. You said a while ago earlier, you did not want one agency to tell another agency what to do.

Dr. NELSON. Yes, sir.

Senator FORD. I wish we could tell that to OMB. [Laughter.]

Senator FORD. You are mighty nice. And we do appreciate your hard work and your cooperation. There may be some questions in writing, and we hope you respond to those in a timely manner.

Dr. NELSON. Thank you, Mr. Chairman.

Senator FORD. Thank you.

And at the request of my good friend Senator Dominici, we will now have a panel consisting of Dr. Rollwagen, Robert Paluck, Dr. Connelley and Dr. Hecker. And if all four of you gentlemen will come forward, we will hear your testimony and then ask some questions.

Dr. Rollwagen, if I have pronounced that right, or just Rollwagen?

Mr. ROLLWAGEN. Rollwagen is fine. The doctor is a promotion.

Senator FORD. Sir? You may proceed. We will work right down the table to Dr. Connolly.

STATEMENT OF JOHN A. ROLLWAGEN, CHAIRMAN AND CHIEF EXECUTIVE OFFICER, CRAY RESEARCH, INC.

Mr. Rollwagen. Thank you very much, Mr. Chairman, for this opportunity to speak before your committee and make it clear at the outset I support this bill, S. 1976. I and my company very much support the bill, its concept, and what it is trying to accomplish. And we think it is critically important for the country as a whole.

I will not go into great detail. I have a prepared statement, which I will submit.

Senator FORD. We would appreciate it.

Mr. ROLLWAGEN. Let me just make a few points though, if I may, that appear in the prepared statement and a couple that have come to me as I have listened to the testimony here.

First of all, I absolutely want to acknowledge the significant part that the Department of Energy and particularly the Los Alamos

laboratories have indeed played in the development of our company Cray Research. I think it is fair to say that Cray Research would not exist today, if it were not for the Department of Energy and Los Alamos Laboratories, that the support has been critical to our success.

This started in our very early days with our first installation, as I think you all know, in the middle 1970s and continues today through a whole range of collaborative efforts that we have worked on together that has produced both hardware and software innovations that have been incorporated in our product and have been made available broadly in the world really of high performance computing. And continues with discussions even today of how we can incorporate developing architectures and technology in our products and services to continue this trend.

We see this bill as a way to continue that collaborative effort and a way to accelerate the process greatly. We also recognize that it is a companion bill and operates as a part of the overall structure developed by Senator Gore and Congressman Walgren last year in their bills on high performance computing. And as Dr. Bromley said, it is consistent with and complimentary to the administration's—at least the principles they are talking about.

And I would like to talk a little bit about that, because what I am hearing so far today is that people's hearts are all in the right places and seem to be in agreement. And frankly, I think it is time we see the feet moving in the same direction too. And I believe that we have made progress in the past with existing structures. But the current times, and particularly, the decade of the 1990s calls for much more than what we have done in the past. And that we need to move forward.

The 1990s in my opinion are a new age in super computing. It is unprecedented; it has not existed before. And the 1990s require new initiatives in super computing. And I believe that the Federal Government has a very major role to play in that regard, particularly, in the area of—if I can put it this way—developing infrastructure which can support the technical developments that are going to be taking place.

I think there are three basic reasons for the fact that the 1990's are a new age, that they are different from the past. First of all—and this was referred to, I think, by Dave Nelson—super computing and the concept of simulation has produced an entirely new third route to new science. The previous two having been in place really for 300 years, ever since Newton and Descartes, being mathematical theory and empirical experimentation.

The third route is numerical simulation, which in a sense combines the other two and makes possible the idea of running simulated experiments that are much faster, much more accurate, much less expensive, much safer than physical experiments. And that those experiments themselves are based on the mathematical—

Senator FORD. When you say safer—I apologize for doing this—but you mean security wise?

Mr. Rollwagen. No, I mean physical safety.

Senator FORD. Oh, okay.

Mr. ROLLWAGEN. For instance, probably these days. the largest single application area for Cray computers, at least new contracts, new orders that we are involved in are in environmental studies rather than, for instance, weapons design. And if we are going to understand and do something about climate change, it is only through exploring the topic with simulation. Because we cannot run physical experiments in the atmosphere. And they take too long, and they could destroy the whole climate. So what I mean is safer in that sense.

So there is a new third route to science that has not been in place before. It is just now coming into place. Secondly, we are taking on much more challenging, significant issues. And these have been summarized in various places. And the best phrase to describe them are called is the grand challenges that we are proposing as a government to take on issues that have not been tackled before, because they just simply could not be.

And these would include, for instance, the greenhouse effect and climate change in a number—and these were referred to in my testimony actually—a number of different scientific areas that require resources, scientific resources well beyond anything than we have had in the past. This is new. This is something we have not done before.

The third thing that has occurred, and that is a specific subject to this bill, is in the realm of networking. The fact is that networking the technology itself has developed to the point now where these large scale systems can in fact be provided to researchers and scientists anywhere in the world, at any time of the day or night and in a very effective way, as if they had the super computer right on their desk.

I would be surprised if as of this point in time more than 10 percent of the scientists and engineers in this country or in the world for that matter have access to or have used super computers; 10 percent would be a big number in my opinion. I do not have the statistics, but that would be my guess.

With networking, it is literally possible to deliver this new capability, this third route to science, to every single scientist and engineer in the country and the ability to have this at their fingertips can make a very fundamental change in how they do their work, and the rapidity with which they do their work, and the effectiveness with which they do their work. They are able to explore 10 times as many alternatives at least, as they would with normal methods.

When you combine these three factors together, to me, it says that we are in a new world that supportive as existing structures have been—and they should be continued—that this is simply not enough. And that there is so much at stake, we must do more than we have in the past. And this is one initiative that can help us do that.

The benefits can come to us in terms of technology transfer among the various laboratories that are work, in the government, in academia, in commercial enterprises. We would literally push the envelope of our own technology in this country making it available to scientists and engineers around the country. We have to recognize that we are not alone on the planet. And that particular-

ly the Japanese have recognized this. And they are already in many areas, particularly, in automotive, chemical, pharmaceutical and electronics areas using this technology today more than we are in this country.

Another factor that is very important though is that it is one thing to have the facilities available, it is another thing to know how to use them. And that the education aspects of this initiative are also critically important as are the software aspects that we have talked about—or that have been talked about earlier today.

And it is not technology just for the grand challenges either. It also turns out to be very important for what would be considered mundane activities. One reference I can make that is in my testimony is to a center for believe it or not advanced cement-based materials located at Northwestern University, who is going to be using 150 hours of Cray Research computer time at the MIT facility to talk about or explore new cement compounds that will affect the structural integrity of that product. So this is applied in many, many areas as well as the grand challenges.

The point is that super computing is beginning to touch all our lives in everything we do and is providing tremendous benefit already in terms of safer cars and less expensive oil—although all this stuff is expensive enough—better weather forecast, safer planes, et cetera, et cetera. Furthermore, our economic leadership is tied to our ability to continue to move forward in these areas.

And I will simply close by choosing from Dr. Bromley's testimony my favorite line, which is to say that there is no—this is the single most important thing that we could do to promote science and technology in this country. It seems to me that is a pretty good reason to go forward.

[The prepared statement of Mr. Rollwagen follows:]

Statement of John A. Rollwagen

Chairman and Chief Executive Officer

Cray Research, Inc.

on

The Department of Energy High-Performance

Computing Act of 1989

S.1976

Before the

Senate Subcommittee on Energy Research and Development

March 6, 1990

Good afternoon. My name is John Rollwagen a I am Chairman and Chief Executive Officer of Cray Research, Inc. Thank you Mr. Chairman for the opportunity to testify before your subcommittee on S.1976 - The Department of Energy High Performance Computing Act of 1989. The DOE has played a seminal role for the supercomputing industry. Laboratory scientists are among the most cycle hungry users and the laboratories have been leaders in demanding leading-edge supercomputing. The DOE, in my view, not only has one of the most aggressive science programs in our country, but it also has one of the most well-defined and institutionalized processes for technology transfer.

Cray Research, Inc. is pleased to have had, from our very beginning, a close collaborative relationship with the DOE laboratories. In 1976 we were a new company and the CRAY-1 was a "bare-bones" machine with no application software or operating system software. We needed a very competent and capable customer to make it a success. We asked the Los Alamos Laboratory to accept our machine on a loan basis and evaluate it. We realized the benefit to us if this evaluation was positive - and it was. In this regard, the DOE has played a special role in history of Cray Research. As examples of specific cases of collaborative efforts between Cray and the DOE laboratories, I have attached a list of products developed at Los Alamos and Lawrence Livermore laboratories directly adopted by Cray Research, Inc. (see attachment 1)

Since 1976, Cray Research has grown to become the world's leader in the supercomputer industry and the issue of high-performance computing has attracted not only the attention of the scientific and engineering community, but also that of U.S. Government policymakers. We are pleased to work with Congress and the Administration which have both recognized the importance of the issue and have initiated programs to address the need for a federally coordinated effort for high-performance computing.

Cray Research supports S.1976. It appears to be an integral part of the overall structure that has been outlined in the programs defined in more expansive legislation introduced last year by Senator Gore (S.1067) and Congressman Walgren (HR 3131), The High-Performance Computing Act of 1989. Also, we believe that S.1976 is supportive of the Administration's Office of Science and Technology (OSTP) report entitled, The Federal High-Performance Computing Program. The DOE High-Performance Computing Act highlights both the leading-edge R&D positioning for the future as well as investments in the infrastructure, to greatly broaden the industries and users of the technology.

Cray Research endorses all these initiatives because they will act as catalysts to encourage government agencies and private companies to join together and fully utilize the enormous benefits that can be realized through high-performance computers.

One of the most critical elements to insure success of the high-performance computing program is the effectiveness by which the fruits of its R&D are rapidly transferred back to the U.S. industrial sector. This is an area of traditional weakness in the U.S., especially when compared to Japan. Improvements in dissemination are impeded through the rigid sector definitions of the U.S.: university, government, private - with different R&D policies for each. Currently, the only way that U.S. industry can take government R&D dollars for advanced R&D projects is to become a government contractor. This is by many accounts a costly and onerous undertaking, and many small and medium-size companies cannot afford to do this.

We believe that S.1976 is a significant step in the formation of cooperative DOE research and development centers and hopefully will help improve technology-transfer mechanisms to ensure that the results of basic research are readily available to U.S. industry. From industry's point of view, the monitoring, interacting and transferring of technology from national labs and universities can be a costly and time-consuming undertaking. Significant computational R&D is undertaken in a plethora of local laboratories as well as some national centers--far too many for individual companies to track. More government coordination of technology transfer issues, would help increase the flow of technology from the labs to the commercial marketplace.

It is our hope that these individual R&D efforts proposed by S.1976 will contribute to the national technology base and that somehow we as major stakeholders can both become aware of these activities and benefit from the results of already funded R&D.

Regarding the technology base issue, I want to point out that Cray Research continues to be a significant investor in supercomputing R&D. For example, in 1989, we invested $11 million in R&D grants to university researchers focussing on leading-edge computational science. Since we initiated this R&D program in 1985, our investments have totalled approximately $30 million. This translates into about 170 on-going projects spread over various schools and institutions. Cray Research also donated $4 million to NSF's Science and Technology Center's Program in 1989, recognizing that the "size" of the science follows the size of the computing tool, and hoping to accelerate the shift to the computational approach to science, particularly in these centers which have been nationally selected as centers of leading-edge science.

In addition to science and university contributions, Cray Research has and continues to make significant contributions to other elements of the supercomputing technology base. This includes several U.S. semiconductor firms with whom we collaborate in an effort to develop advanced chips that can be used by other computer manufacturers. I mention these examples

of Cray Research's contributions to illustrate how we as a company are investing in this technology base that underpins U.S. supercomputers leadership. We feel strongly that in addition to exploratory R&D, there remains a strong need to focus on R&D technology issues critical to growing a strong domestic supercomputer manufacturing base as well.

This legislation, like S.1067 and HR 3131 helps to foster collaboration among the various governmental agencies involved with high-performance computing by establishing a national multi-gigabit research and education computer network. All sectors of the U.S. infrastructure must work together to meet the challenges that lie ahead. While the U.S. has the largest R&D expenditure in the world--approximately $132B--we work hard to separate our policies and funding activities based upon basic or applied research, advanced development, etc. Even further, we have separated our policies and our R&D support for the university, private and defense sectors. The result is that only a fragment of relevant R&D finds its way into the commercial base to help our overall competitive position.

S.1976 also provides for advanced software technology support, needed for research groups to collaborate to address the "Grand Challenge" problems, such as global climate change, superconductivity, human genome, nuclear fusion; or other large government-sponsored R&D programs, such as National Aerospace

Plane, Space Station, DOE's Environmental Restoration. R&D Megaprojects such as these, characterized by significant science components, are believed to be beyond the capabilities of today's supercomputers. Such projects push the technological envelope for the entire supercomputing base. An important driving force for us and the supercomputer industry is the quest to solve these leading-edge problems of science and engineering. Besides the inherent merit and benefits these projects embody, they will also push supercomputer researchers to develop new software and solution techniques to meet the challenges presented by the problems. The computational solutions to these leading-edge research problems complement the traditional approaches of experimentation and theory, and provides basic understanding to complex problems. It is the application of the supercomputer, the resulting codes and products of their use that provide economic benefits beyond those the supercomputer industry itself. It is understood that economic leadership goes to the country that not only has the best tools but uses these tools for economic benefit. This point can be illustrated on attachment #2 which shows the comparison between the U.S. and Japan in the use of supercomputer technology. The concentration of supercomputer use in Japan is in industries considered vital for 90's - i.e., electronics, pharmaceutical/chemical/biotechnology, and construction.

Economic leadership and globally comparative advantage is increasingly dependent on technclogical innovation and science-based R&D. Many times developments in various industries are limited by the current state-of-the-art of the science disciplines which underpin them. It is my belief that all industry can be helped by the more rapid adoption of new technology. Even a basic industry such as the cement business can benefit from supercomputing. The Center for Advanced Cement-Based Materials at Northwestern University is a NSF Science and Technology Center that has been granted 150 hours of Cray Research time at MIT. The group will use the Cray Research time to study the structural integrity of new cement compounds, different chemical composites for new cements, fluid flows and other applications. This shows how new breakthroughs have been made possible in a traditionally low-tech industry. The improved understanding afforded by the computational approach can lead to lower costs and faster innovation - essential in the global marketplace of today. I believe it was management leader Peter Drucker who said (if I might paraphrase) that companies should take the initiative to obsolete their own technology base or their competitors will do it for them. This notion is equally as compelling in the global arena.

In summation, our message here today is very simple. Supercomputing remains today an innovation that touches us every day and improves our lives now and in the future. Our cars are

better and safer, our oil is less expensive to find, our weather forecasts are better, our airplanes are safer and more fuel efficient. But we are just beginning to make progress in computational chemistry and biotechnology areas. And we are just beginning to understand the global climate change of our planet and how we can affect it. By fortuitous timing, supercomputers are available today to test the complex hypotheses describing the "Greenhouse Effect." For example, a simple 30 year climate simulation testing global warming as a result of CO_2 increases requires about 500 hours to run on a CRAY X-MP/48. Without this computational resource, these simulations would not be possible. Preventive measures will be based on these simulations to minimize the soci-economic impact on our nation economies as a result of climate change. This legislation supports these initiatives and will help this nation achieve greater economic benefit.

I would like to close with a quote (with a slight paraphrase) from the White House Science Council Committee on Research in Very High-Performance Computing:

> "The bottom line is that any country which seeks to control its future must effectively exploit...(supercomputing). A country aspiring to military leadership must dominate if not control supercomputing. A country

seeking economic strength in the information age must lead in the development and application of supercomputing in industry and research."

Thank you again for this opportunity to testify. I would be happy to respond to any questions.

ATTACHMENT #1

Examples of products developed at Los Alamos and directly adopted by Cray Research, Inc.

HARDWARE

Cray Research implemented several hardware features in its product line that came about from collaboration with Los Alamos.

- Los Alamos discovered the "vector recursion" feature with Serial Number 1 (S/N 1) in 1977/1988. Initially, Los Alamos considered it a design idiosyncrasy, but the Laboratory realized its potential and took advantage of it in the software.

  Cray Research adopted the idea to exploit this hardware and modified the Fortran compiler accordingly. The "vector collapse" allowed a user to quickly sum elements in a single vector utilizing the vector recursion.

  Newer Cray Research models no longer contain this hardware feature.

- In another experience with S/N 1, Los Alamos found the Cray Research memory could be enhanced with error checking in the hardware. Although the design did not originate at Los Alamos, the SECDED error checking (based on the Hamming code) in Cray Research computers is a direct result of the Los Alamos influence.

- Sometime in the early 1980s, Los Alamos conceived the notion of a hardware "gather/scatter" capability for indirect array referencing in memory. The Laboratory approached Cray Research with the idea, and other Cray Research customers independently expressed an interest in that capability. Cray Research developed a CPU functional unit as a result.

- The design of the new Input/Output Subsystem (IOS) Model E was influenced by Los Alamos. The new IOS will be available with new systems in 1991, and it will contain redundant paths and error checking for enhanced reliability. It will also be programmed in a higher-level language for versatility and maintainability.

- Los Alamos developed the concept of a standard high-speed channel for different vendor computers to communicate with each other at faster rates. Using

Cray Research high-speed channel speeds as a model, the Laboratory defined the HPPI standard which Cray Research and other vendors will offer in future product lines.

SOFTWARE

Cray Research adopted some software from the National Laboratories for incorporation into released Cray Research products. They are utilities developed at Los Alamos.

- the TEDI editor. This is a "line" editor used to modify source code computer programs.

- the Cray Research SEGLDR came from the CTSS LDR developed at Los Alamos and Livermore National Laboratories.

Cray Research also modified key software to meet the needs of the Laboratory. In doing so, Cray Research and its customers have benefited from the ideas originating at LANL:

o The Cray Research Fortran compilers are modified to use two important ideas developed at Los Alamos. Those are partial optimization to enhance debugging, and Fortran pointers to dynamically manage memory.

o The new Cray Research debugger (CDBX) includes features recommended by Los Alamos and adapted from the LANL DDT debugger.

o The Cray Research runtime scientific library (SCILIB) incorporates ideas and actual code originated at Los Alamos. These are for fast mathematical procedures commonly used by many Cray Research customer codes.

o The new Cray Research operating system (UNICOS) has many features designed by Cray Research to meet the needs of Los Alamos. Cray Research gets the benefit of the Laboratory's 20 years of interactive production computing, and all Cray Research UNICOS customers realize the same benefits. Some examples of the features influenced by Los Alamos are: multilevel security, load leveling and batch processing, resource allocation, and interprocess communication.

ATTACHMENT #2

RELATIVE ESTIMATE U.S./JAPAN ADOPTION OF
COMPUTATIONAL APPROACH TO R&D
(number of supercomputers)

	UNITED STATES		JAPAN	
COMMERCIAL APPLICATIONS				
Aerospace	16		0	
Automotive	6		12	
Chem/Phar/Biotech	4		8	
Construction	0		11	
Electronics	2		14	
Energy	2		4	
Petroleum	19		0	
Service Bureau	3		10	
Finance	0		2	
Other	0		1	
Subtotals		52		62
GOVERNMENT APPLICATIONS	90		12	
UNIVERSITY APPLICATIONS	31		30	
TOTALS		173		104

Source: Cray Research, Inc. estimates 1/90.
(Does not include IBM installations.)

Senator FORD. Thank you very much. Mr. Paluck.

STATEMENT OF ROBERT J. PALUCK, CHAIRMAN, PRESIDENT AND CHIEF EXECUTIVE OFFICER, CONVEX COMPUTER CORP.

Mr. PALUCK. Good afternoon, Mr. Chairman, and members of the committee. Thank you very much for inviting me here today. I appreciate the opportunity to comment on Senator Johnston's bill.

I enthusiastically and strongly am in favor of the bill, because I believe it can be used, if properly implemented to address America's largest problem right now. And that is, our loss of domestic and international competitiveness amongst our American companies and products.

I think in addition it will allow us to stimulate and encourage the development of a very strong super computing industry, as well as being able to accomplish some Department of Energy research and education.

Before I begin my comments on the—my feelings and issues on the bill—I would like to point something out. And that is Cray certainly is the largest most successful and consequently one of the most well known super computing companies in the world and certainly in America. Let me point out that there is another American success story and that is CONVEX. And we have not been that well known here in Congress, and so if I can just spend about 60 seconds to let you know what we have accomplished.

Senator FORD. Take 90.

Mr. PALUCK. Pardon me?

Senator FORD. Take 90.

Mr. PALUCK. We started our company about eight years ago. It started by two people, myself and Steve Wallach a computer architect. It started in Dallas with about $5 million of venture capital money. In the last five years that we have actually been shipping product, we have now shipped very close to $.5 billion worth of product to over 33 countries around the world.

Essentially, we started the company with this recognition back in 1982. We saw that scientists and engineers around the world really needed to have access to super computing to develop new products. The super computers that were in existence in those days were very, very expensive. And because of that expense, it limited the access to those type of machines.

So what we recognized is that what the world really needed was simply a machine that was affordable. A machine that could be easily deployed throughout industry and throughout research. And so what basically we developed as a product line was super computing for the masses.

What has evolved now that we have a first generation product and a second generation product is actually the definition of super computing has changed. And in fact, what has evolved is the same thing that has happened in mainframes, in work stations, in mini computers. And that is that there is a low end, a medium end and a high end.

Cray without any question dominates the high end. We without any question—I think all the statistics show—dominate the mid-range. We actually have products in the low, the medium and the

high. But our principal area of dominance is in the mid-range. Both companies contribute to what is going on in the world in terms of super computing.

Within the last five years we have now sold over 600 systems. That represents about 400 different customers around the world in 33 different countries. About 45 percent of our business is here in the United States, 35 percent is in Europe, but 15 percent is in Japan. We have had very successful activities there.

This last year we did about 160 million in sales with about 12 million in profit. That represents about 50 percent growth in revenue and a 100 percent growth in profit. And this year we are on the New York Stock Exchange, and analysts are predicting we will do about $.25 billion this year in sales. So we are approximately at this point in time pretty close to about one-third the size of Cray. But we are also growing at about 50 percent per year.

As I had mentioned, we have been successful in Japan. And I just returned from Japan on Sunday. And one of the things that really struck me while I was there is how the economy is absolutely booming. The Japanese have a lot of yen that they are spending. And you know a lot of books have been written on why, why have they been successful. But the one thing that strikes me, the one underlying accomplishment that they did is in a broad area of industries and a broad area of markets.

What they did was develop outstanding, internationally competitive products. And that has fueled the economy to allow an upward spiral to occur, where success has kept feeding on success.

One of the reasons why I am enthusiastic about Senator Johnston's bill is that I believe if properly deployed, it could do the same sort of thing. It could be a stimulant to help American competitiveness with regard to the products that are being generated.

Super computers, of course, are a basic building block. They are like the microscope to a biologist. They are like a telescope to an astronomer. They are tools. They are fundamental tools. But put in the hands of the biologist and astronomer and the car manufacturers and the petroleum industry and on and on, innovative products can be brought to market. And they can be brought to market much quicker than the competition.

And so the fundamentals of the telescopes and microscopes or super computers is interesting. But it is really the ramifications of the accessibility of those tools that can help change the overall productivity and the economy in the United States.

So with that, I would like to explain now the four thoughts with regard to the bill. The first is throughout the bill there is a discussion of partnership that the Department of Energy needs to work in partnership with industry. I want to underline that and highlight it, because I believe that the Department of Energy has phenomenal resources in the area of research.

Industry has some phenomenal contacts with regard to what is actually going out in the market, and with regard to how other countries, how other companies, how other research institutions actually deploy these tools, actually use these tools. With the proper marriage, with the proper cooperation and partnership the money can be spent, each dollar that is spent could be directed with the

bias towards increasing the overall productivity of the United States, but especially the competitiveness of products.

There is a risk that is always run in a program like this, that money is used for research for research's sake, which is good. But the threat that we have to fight today is the increase in the productivity and the competitiveness of American products. That will solve the problem. And so if the bias of the way the money is spent is towards that, and we have made a very broad improvement on what is going on with regard to the success of American, not just success of super computing.

The second point that I would like to make is that I believe that in the implementation of the network, a network that uses a multi-tired type approach should be used. And let me explain what I mean by that. Essentially there is traditionally—especially years ago—two different types of, let me call them, religions in terms of—or philosophies—in which to implement a system of computers. One is the centralized approach, where everyone that uses that computer essentially shares that one resource. Another religion is a called a distributed approach where a department or a small group of people have their own. And they have their own control over it.

What is evolving today is a major trend in industry and in government agencies as well. And that is to implement a multi-tired approach where departments and activities and projects have their own super computer to be able to have their own control, to be able to run large applications, set their own priorities, control their own destiny. But for the largest jobs, the largest jobs do require the fastest machines in the world, which are also the most expensive machines in the world which should be centralized.

And so what I am proposing is that a multi-tired approach be implemented to be able to satisfy the needs of the user. A good example of this is the very successful National Science Foundation's super computing centers. Certainly it has been a major contribution, I think to research, in some degrees to industry. They asked the users what they would like; they would like to have both, because too many problems that are actually too small to be used on the large machines are clogging up the large machines. And so a combination of both ends up being a very efficient way to accomplish things.

The third point I would like to make is to highly encourage and highlight one of the aspects of the bill that talks about buying the first production or the original prototypes, from super computing companies. I think this will help in two ways.

It certainly helps the development of new programs, because there is confidence there that there are people to buy it. I also comment that in the development of machines that are, let us say mid-range and not the fastest machines in the world, they also require very sophisticated technology to be able to achieve that affordability. And so this will really help, it will really help the industry in general. But it has got a second, it has got a second real benefit I think to American super computing manufacturers, certainly for the Department of Energy.

Companies, other companies other than Cray, lust for the fantastic relationship that has evolved between Cray and the Department

of Energy. It exists because Cray has done such a phenomenal job of products, of services, of support and so on. And so that relationship exists, but it does make it somewhat difficult to break into to get in the door, for other companies to be able to have that relationship also. And this bill—and this aspect of this bill—will actually help to open up the door.

I heard the comments just moments ago about how other companies are in there. That is true. CONVEX also by the way has a machine at the Department of Energy—we are talking about quantity. Other agencies, for instance—other departments of the government have dozens of machines. Whereas in the examples that might have been given before, we are talking probably about one machine or two machines. So anyway, this is very healthy. I think it is very healthy for everyone to allow that to happen.

The fourth area is that I believe that to make another real contribution to the super computing industry, which then can benefit the overall business world is that projects—that commercially developed projects could be funded by this particular bill. And let me explain why.

American companies have a real problem in terms of finding the necessary profit and loss balance sheet dollars to be able to develop new products. Let me explain why. Companies like CONVEX for instance have a lot cash, Cray has a lot of cash, other companies have a lot of cash. But we are limited by what the investment community in the United States will allow essentially us to spend from a profit and loss standpoint.

With the context that we have in the market, the opportunities that we see in the market, there are some phenomenal innovations that can be done, phenomenal additional products that can be developed that will benefit Department of Energy computing and simultaneously be commercial products that can be used by industry in general. And the ability to have essentially off balance sheet development would be a real edge. And in fact the type of edge that we need to be able to compete with countries like Japan, where they have other advantages, which I would not even call unfair. They are just advantages. And this could give us an advantage.

So in summary, I strongly endorse the overall plan. I believe that there are four areas that should be highlighted: the partnership, the network multi-tired aspect, the purchase of advanced machines as well as the funding of commercial developments in conjunction with the Department of Energy.

[The prepared statement of Mr. Paluck follows:]

28-910 O - 90 - 4

Statement of Robert J. Paluck
Chairman, President and
Chief Executive Officer
CONVEX Computer Corporation

on the

Department of Energy High-Performance Computing Act of 1989

before the

Senate Subcommittee on Energy Research and Development

of the

Senate Committee on Energy and Natural Resources

March 6, 1990

Robert J. Paluck

Good afternoon, Mr. Chairman and Members of the Committee. I appreciate Senator Ford's invitation to be with you this afternoon to comment on Senator Johnston's Department of Energy High-Performance Computing Act of 1989. Also, since you may not be familiar with CONVEX Computer Corporation, I would like to tell you about our success in supercomputing and the strategic role we play in high performance computing around the world.

Introduction

I am pleased to testify today on this Bill, because I strongly believe it can positively contribute to one of the biggest threats facing the United States today: the loss of our international, as well as our domestic, competitiveness in all technology related businesses. The Bill will, in addition, directly stimulate the American supercomputing industry, as well as further the research objectives of the Department of Energy. Its effects are pervasive and truly a WIN-WIN situation for the United States.

Supercomputers are the fundamental building block tools that contribute in almost all disciplines across the broadest spectrum of science and technology. In the 1990's, the way America can stay competitive is literally to put supercomputing in the hands of the "masses." Supercomputers are to the modern technologist what the invention of the microscope was to biologists and the telescope was to astronomers. In fact, supercomputers enable scientists and engineers to solve problems for things that are too small, too large, too quick, too slow, or too dangerous to observe directly. Their use in industry results in new products that are more innovative, safer, and get to market quicker. Their use in research results in fundamental breakthroughs in sciences that change how we see the world. The supercomputer is the one common tool across all U. S. scientific and technological activities that, if put in the hands of engineers and scientists throughout the United States, can dramatically sharpen the competitive output of the United States.

Of course, Japanese industry and research institutions totally understand and believe these concepts. From our perspective, they have been the fastest nation to purchase CONVEX's latest technology. Until just recently, there were more of CONVEX's top-of-the-line supercomputers in Japan than in the United States. American researchers and engineers believe these concepts also, but access to

supercomputer tools has been limited. This Bill, along with Senator Gore's S. 1067 Bill, can be the catalyst to change this trend.

Since you may not be familiar with CONVEX Computer Corporation, I would like to first tell you what this American supercomputer company has accomplished over the last eight years. Then I want to provide you with some perspectives in high performance computing from CONVEX's viewpoint, using these remarks as a framework from which to comment on the Bill.

CONVEX Computer Corporation

CONVEX Computer Corporation is the leading supplier of mid-range supercomputers in the world. Steve Wallach and I founded the company in 1982 with $5 million of venture capital money. Now, after just five years of shipping, we have shipped nearly a half a billion dollars of systems to customers in 33 countries around the world, and have been growing at 50% per year.

We started the company based on the recognition that scientists and engineers needed supercomputers to get their jobs done, but few could afford the initial investment as well as the on-going support costs. So, using advanced semiconductor, software, and architectural technologies, we developed the first affordable supercomputer -- at that time called the first mini-supercomputer. Our first machine had the near performance of a Cray 1, but at a fraction of the cost. In 1988, we began shipping our second generation products, which is a family of seven supercomputers that range in performance from our original product up to our high end which is as fast as a Cray X-MP/14.

Technology organizations today view high performance computing in a three-tier hierarchy model. The bottom tier is composed of workstations and terminals that are excellent at being the user interface and for running small applications. The middle tier is the high performance machines that the workstations and terminals are connected to, but can also be owned, located, and controlled in the department or assigned to a project. This is where mini-supercomputers have contributed. These machines can also solve very large problems, plus they can be distributed close to the user. The top tier is the machines that must be centralized and carefully shared because of the very high cost, and should be used only for the

applications that need the very highest performing machines. In this way, the size of the application is matched to the most appropriate resource and, consequently, an efficient, cost-effective, overall network of systems is achieved, as well as satisfying the needs of the users.

CONVEX's success has been based on providing systems for the middle tier and, recently with the availability of our fastest system, the top, centralized tier. We now have shipped over 600 supercomputers to 375 customers in 33 countries, including nearly 100 systems to Japan. Our machines are being used at Ford and General Motors to design more competitive automobiles; by Mobil, Texaco and Shell to find oil; distributed throughout NASA and the Department of Defense for a wide range of design programs; as well as the National Institutes of Health for the genome project, and over 120 universities for every area of science and engineering.

Because of the size, performance and affordability, we have more systems installed in American government, industry, research and universities than any other supercomputer manufacturer. Globally, we have shipped to more countries than anyone else. Whereas Cray Research dominates the industry with the highest performance machines, CONVEX dominates the industry at the mid-range. Both are needed to achieve the most cost-effective overall system, as well as to satisfy user needs.

CONVEX has had 21 consecutive quarters, that is five full years, of increasing revenue and has been profitable since just after the first product shipments. Forty-five percent of our revenue comes from the United States, 35% from Europe, 15% from Japan, and 5% from other countries, such as Australia, Brazil and Taiwan. Our stock is traded on the New York Stock Exchange, and securities analysts are predicting we will ship a quarter billion dollars of systems in 1990.

As you can tell from this brief story, we are a successful American supercomputer company. We are growing very fast and being exposed to how our customers are benefitting from our systems all over the world. It's this exposure that makes it so clear to me how the overall competitiveness of the United States, as well as the American supercomputing industry, can benefit from the proposed Bill, if it is properly funded and enacted now.

Robert J. Paluck

I would like to share with you some perspectives on supercomputing that can help put the Bill in a useful framework.

CONVEX Perspectives

The supercomputing market today is a global, international marketplace dominated at the high end by Cray Research, and dominated at the mid-range by CONVEX. The supercomputing market has evolved over the past six years from being narrowly defined as a point product, i.e. "the fastest machine in the world," to a broad choice of products that offer a range of performance at various prices. This is no different from what has happened in every other computer market; there is a low end, a medium range and a high end. All three have proven to be needed to meet the needs of industry and research. This is proven by the number of our systems sold throughout the supercomputing industry.

The Japanese computer manufacturers have high performance, extremely fast hardware products which they have brought, or are about to bring, to the high end of the market. However, at the present time, the United States high performance computer companies continue to lead the world for one prime reason: supercomputing is not merely about the fastest hardware; it is all about solving complex problems. American companies remain ahead of foreign competition, because we understand that there are four key aspects to supercomputing solutions.

First, there is balanced, high-performance hardware. There is more to real computer performance than pure megaflops or gigaflops. Unfortunately, that is how performance is commonly measured, but these measurements must be properly interpreted. Useful computer performance is much more than peak performance; it is software performance, memory performance, and I/O performance, as well. Supercomputer users care only about the performance of their particular applications, i.e. the problems they solve with their machines. Success in this type of performance is determined by dozens of attributes. In terms of pure, native speed of the hardware, the Japanese are challenging the United States strongly at the high end, but that is only part of the supercomputing success equation.

When computers were first developed, they were employed only as central site facilities serving entire organizations. There were many valid reasons for this, not the least of which was cost. When

minicomputers came along, the notion of distributed data processing came into being. Providing very useful, but not the most powerful, computing close to users was judged to be a wise use of this less expensive technology. It was done, not to the exclusion of central site facilities, but in a complementary manner. We now have the exact same situation in the supercomputing industry. Cray dominates the high end, because it provides the fastest machines in the industry. CONVEX fills that complementary role, providing the most innovative, leading edge technology in its price/performance class. There is nothing contradictory in our corporate philosophies, or the strategies that we have individually adopted in pursuing our separate, complementary portions of the supercomputing market.

Second is system software technology. In this aspect, I am primarily speaking of operating systems and compilers. While American technology has always been on the leading edge in system software, we have historically made this one of the strongest bastions of proprietary stubbornness. By that I mean that each company had its own operating system software, closely tied to its own hardware, and largely incompatible with the operating systems of other manufacturers. We have now entered a new age, with CONVEX, Cray and many others embracing standardization with the UNIX™ operating system. Not only has CONVEX embraced UNIX™ as our operating system, but we have designed our hardware with UNIX™ in mind. Compilers, the parts of the system software that translate user-written programs into machine language, are currently produced best in the United States. The Japanese lag behind, but not by much.

The third important aspect is applications software. Most of the supercomputers in use today, and CONVEX is no exception, utilize third-party written applications software, rather than custom-written programs. The greatest majority of these third-party software packages is developed and sold by United States companies. CONVEX considers having a broad array of applications software available on its machines, and having close business relationships and agreements with these software companies, as critical to its current and future competitive strategy and success. The true size of the market that a computer company has access to is set by the number and variety of third-party software packages that are offered for its machines. As of this moment, there are over 425 separate and distinct third-party, UNIX™-based applications software packages available for use on CONVEX machines. This is

the highest number in the industry. American companies continue to be the premier software development houses in the world.

The final aspect is service and support. Simply taking care of the customer is what I am addressing here. This is a critical, if not *the* critical, component in supercomputing solutions. Users of supercomputers, not manufacturers of supercomputers, are working diligently on solving some of the world's toughest problems. Supercomputers are their most important tools, but tools they are. These sophisticated problem solvers require all of the service and support that we can give them in their endeavors.

Currently, American supercomputer manufacturers lead in all four areas, but that leadership can be a fragile and fleeting thing. Our Japanese supercomputer competitors are well-financed, highly integrated companies with very long-term technology investment perspectives, that the American equity markets do not permit us to have. They clearly have as their goal, and have the capabilities, to build the fastest hardware in the world. They have an impressive track record of incessant improvements to their products that can bring their systems software to par with U. S. technology. Sales success, especially in Europe and now Eastern Bloc countries, will motivate third-party applications vendors to port their software to the Japanese platforms. And even with the best service and relationships, history has shown in other markets that consumers all over the world will buy from companies that offer the best total, overall corporate products. So American leadership is clearly threatened and action is needed now!

"Success fosters more success," applies to Japanese computer manufacturers, as well as American. If American manufacturers are delayed in getting export licenses, the business will go to the Japanese. This is especially true for the opening up of the eastern European countries. Certainly U. S. defense is a most critical issue, but technology is advancing so rapidly that equipment designed just five years ago is now obsolete by Government procurement standards, yet not exportable. Export thresholds must be practical, and revised commensurate with advancing technology.

Procurement practices must be fair both in foreign countries and our own. Much has been written about Japan, but what about us? Not all problems need the fastest, most expensive machines to do the job. Federal budget austerity should demand that a range of

supercomputer solutions be considered in every procurement. Sole source procurements can help one American supercomputer company, but at the expense of all other American companies.

Staying competitive in today's global supercomputer market will take a concerted effort by American companies, as well as cooperation and constructive stimulation by Government. Certainly, the DoE High-Performance Computing Act of 1989, when enacted, will be a positive contribution in this direction.

Comments on the Bill, S. 1976

General Comments

I enthusiastically endorse Senator Johnston's Bill, and commend it to you for your favorable consideration and swift passage in the Senate. There are a few provisions of the Bill that I would like to comment on and highlight.

The Federal Government, in particular the Department of Energy, has played a prominent role in the American supercomputing success story. This Bill again demonstrates this leadership. In several areas of the Bill, it is advocated that the Department of Energy work as partners with industry to develop the strategies and, in certain areas, the implementation of the plan. I want to highlight this concept as being extremely important to achieving the objectives of the Bill. I feel that to have the most effect on the competitiveness of the United States as a whole, the results of the expenditures for equipment, and the results of the expenditures for research, must ultimately and eventually lead to competitively superior commercial products. The strategic plan that is put in place by this Bill should have this as a fundamental objective. The Government is better qualified for some aspects of the task, and industry is better qualified for others. Partnership allows each of the partners to take full advantage of the best capabilities of both. CONVEX has exposure to applications, research, and product developments occurring all over the world, and in the broadest of scientific areas. We volunteer to help in whatever ways we can.

I fully support the notion of a National High-Performance Computing Program with OSTP's Federal High Performance Computing Program providing the framework. We support enactment of such a program via such proposals as we are considering here today and those being

considered at other times, such as Senator Gore's Bill, S. 1067. What we and the Federal Government need is a coordinated, effective program efficiently managed so that we all benefit. I see those thoughts in Senator Johnston's Bill and I applaud them.

Let me now turn to some particular aspects of the Bill.

The National Research and Education Network

Supercomputing can be configured or approached in a number of ways. The CONVEX perspective is that too much emphasis has been placed in the past on providing limited access to too few centralized machines. Supercomputing must be made available to, and meet the needs of, a broad base of users through widely distributed supercomputer systems, placed closer to the ultimate user; not supplanting the centralized machines, but rather complementing them.

If we look at the way telephone companies do business, we will discover a very valuable insight by analogy to what I am suggesting to you on this point. Telephone companies have historically provided a hierarchy of services implemented through a hierarchical system of telephone switches. None of us has their home telephone instrument directly connected to the largest of AT&T's toll centers. Rather, we are connected to a much smaller local exchange that adequately provides for most of our needs, because we mostly make short distance, voice calls that do not require the availability of all of the capabilities of larger telephone switching centers. However, when we need those kinds of increased services, such as long distance voice or high-speed data communications services, they are available via that network to our homes and businesses.

What I am suggesting is that we not zero-in on the National Research and Education Network (NREN) as only a multi-gigabit per second backbone network, connecting only a small number of very high-speed, centralized computer systems. Let's think of it as a distributed network of computing and telecommunications services, serving the widest possible number of scientists and engineers from government, industry and academia. The National Science Foundation's national supercomputer centers represent a case in point. The program has been a success, but we can learn from what those users are additionally asking for: supercomputing close to the user. Let's supplement and complement the national supercomputer

centers with affordable, open, accessible supercomputing facilities, available in departments and dedicated to projects, across the Nation. Let's put a broad range of supercomputers, distributed data bases, and other research and production facilities in the very laps of those that need them to help maintain and regain America's preeminence in many disciplines.

Software

Let me move on to software. In the last ten years, only about 300 high-end supercomputers have been sold by American companies to industry and to research institutions. From CONVEX alone, over 600 high-performance computing systems have been shipped in only five years. American industry needs distributed, affordable supercomputing power to remain competitive. These companies, large and small, are voting with their checkbooks for this way of providing supercomputing. They are using supercomputing in production environments, not just in their research laboratories. They need supercomputers to bring new and improved products to market faster. Supercomputers are a valued competitive weapon for all of these companies.

The full utility of supercomputers can only be reached through software. The sophisticated supercomputing user community desperately needs improved software development tools, computer-assisted software engineering (CASE) capabilities, and better algorithmic methods. With this improved state-of-the-art software, we can move forward with attacks on the Grand Challenges, which are enumerated in Dr. Bromley's Federal High Performance Computing Program document. Those Grand Challenges singled out in the Bill are certainly well-suited for solution by the Department of Energy and its world-renowned National Laboratories.

I wholeheartedly support the software tasks and goals of the Bill. Care should be taken to be sure that resources are not wasted by reinventing what may already exist in industry or somewhere in the world. Let's concentrate on improving software technology, but adhering to industry standards wherever possible, and avoiding proliferating proprietary solutions to software problems.

Robert J. Paluck

High-Performance Computer Systems

The industry can certainly benefit from the assistance of the National Laboratories, and their exceptional corps of outstanding scientists and researchers, in performing the necessary research and development of hardware components and architectures, as detailed in the Bill. Again, let's be partners here and use the DoE facilities and people to best advantage.

CONVEX welcomes the opportunity offered by the bill's provision to have the DoE "purchase early market and production model computer systems and subsystems." Currently, CONVEX has three major advanced programs under way that, when completed, will benefit the DoE high performance computing effort. The encouragement of this Bill to buy first shipments is a healthy and constructive way to help American supercomputing.

Basic Research and Education

I strongly support the provisions of the Bill in the areas of basic research and education. In basic research, the National Laboratories have the talent and facilities to supplement the mostly applied research found in industry. Only the largest and richest corporations can afford to have very much of their resources dedicated to basic research. Most of the industry, and I count CONVEX in this group, must use its limited research and development resources in the development and production of the next generations of our commercial products. So we need a fertile source of basic research, if the supercomputer industry and the Nation are to progress.

Again, I stress that we must treat this as a partnership. Supporting basic research by itself is like cultivating fruitless trees. You can enjoy the flowers, but that is all that you get to do. So we must be sure to create effective, efficient, fast-acting technology transfer mechanisms, so that our basic research tree has fruit as well as flowers. I would recommend that the Bill specifically direct the Secretary of Energy to establish and implement a separate, responsible Technology Transfer Program Office to insure that basic research becomes products to be used to further all of our goals.

In the realm of education, the United States needs a lot of help to keep us competitive. The Bill's provisions to educate and train more researchers; to support graduate students, fellows and faculty; to

share National Laboratory facilities; and to give high school students an exciting look at high technology, are all extremely worthwhile. The industry and the Nation need to educate more computer scientists, computational scientists from various disciplines, electrical engineers and materials scientists, as called for in the Bill. However, mechanical engineers, packaging engineers, chemical engineers and others should also be included. The spirit and intent of the Bill should be applied across the board in the supercomputing industry.

Summary

In summary, I recommend this Bill to you. The $675 million over the five year period will provide substantial stimulation for the industry, benefit the research needs of the Department of Energy, and help sharpen our Nation's industrial competitive edge.

Mr. Chairman, this concludes my prepared remarks. Thank you for inviting me here today. I would be happy to respond to any questions.

Senator FORD. Thank you.
Dr. Hecker.

STATEMENT OF DR. SIEGFRIED S. HECKER, DIRECTOR, LOS ALAMOS NATIONAL LABORATORY

Dr. HECKER. Thank you. I was initially to go last. Would you like me to go next?

Senator FORD. Well, us Kentuckians are hospitable, and if you need a little libation to get you through it we could probably furnish that, too, but you may proceed.

Dr. HECKER. I thank you for your hospitality, and I am happy to be here today. I have a detailed witness statement.

Senator FORD. It will be included in the record, Dr. Hecker.

Dr. HECKER. Thank you, Mr. Chairman. The distinguished witnesses that have already gone before me have spoken about the importance and also the timeliness of the high-performance computing program, and so I will indeed not elaborate on that, but I do want to point out that I certainly support the intent of Senate Bill 1976, and in fact the entire computing initiative as outlined in Dr. Bromley's program.

I would like to leave three important messages with you today. The first one is that these advances that we are talking about—the high-performance computing—should be driven by applications. I think that is absolutely crucial. For a computer to be a useful tool, you have got to do something with it.

Second, I feel very strongly that the Department of Energy should play a leading role, and third, I would like to encourage even stronger interactions between the computer suppliers and the users, and in fact I will develop my proposal for what I call collaborative R&D centers to strengthen that relationship.

But before commenting specifically on the bill, what I would like to do is to spend a few minutes to develop the role that Los Alamos has played in the supercomputing business in order to provide the rationale for the recommendations that I would like to make.

Supercomputing originated with the Los Alamos problem. That is what developing, designing and computing for the first atomic bomb was called during the Manhattan Project, and we actually designed and built one of the world's first large-scale electronic computers—the so-called MANIAC—in the early 1950's.

Since then, we have worked with all of the leading supercomputer manufacturers. Initially with IBM, then with Control Data Corporation, followed, of course, with Cray Research, and also recently we have worked very close with Thinking Machines Corporation.

Our role has been principally that of a very demanding sophisticated user with our goal being to maximize the productivity of the individual researcher, and in that capacity we have often developed new software, we have developed new algorithms to be able to use machines and solve problems, and occasionally we have actually developed new hardware.

In fact, in 1976, the Cray story with us began when Seymour Cray came to Los Alamos and offered us the first machine—the first Cray supercomputer. He came without software, without an

operating system, and asked us to help him put this machine on line, and indeed we did, and I think I can say nothing stronger than the comments that Mr. Rollwagen has already said, and that is that the DOE and Los Alamos were absolutely crucial, and Cray would not exist without them.

Today, Los Alamos is the most powerful scientific computing center in the world. Of course, our main workhorses are the Cray machines. We currently have five, including two of the newest Y-MP's. That is the equivalent of 65 of the original Cray 1 computers. In addition to that, we actually have two of the latest connection machines—the massively parallel machines that are manufactured by Thinking Machines, Incorporated. With that center, we serve 8,000 users Nationwide.

Let me turn for a few minutes to the future computing environment as I see it. I agree, again, with the comments of the witnesses earlier that there is no end in sight for this revolution that we call the computer age. Better performance will enable the scientists and engineers to solve ever more complex problems, and I think high-performance computing has yet to make its greatest contributions to business, to industry and in fact to our every-day lives.

However, to be able to achieve this impact and this continued revolution, we must do two things specifically. First, the United States must stay at the leading edge of computing, and second, as has already been pointed out earlier, we must provide supercomputing power to a greater fraction of the population and as Mr. Pollock said, to empower the masses, so to speak, is absolutely crucial if we are going to take maximum advantage of what supercomputers can provide for us.

By empowering the masses, the sort of things that we would see is that we would not only see the increasing role of supercomputers, of high-performance computing in business and industry, but also in the home. I think it is not too futuristic to dream of the family television equipped with high definition capability, with supercomputing power as the brain center of what we could call a multi-media business and entertainment center that enable you to draw upon information that is stored or generated around the globe, and let us say, all at the touch of a mouse at your desk-top.

So what I would like to see is, today, my eight-year-old daughter has command of her Macintosh with a touch of that mouse. By the time she gets to high school, I would certainly like her to have the supercomputing power available at that touch of the same mouse with the same sort of user-friendly interface.

But let me get back to staying on the leading edge, because I think that is where the principal contribution of Senate Bill 1976 comes in from the standpoint of the Department of Energy and the laboratory's contributions.

We need to push several current trends: faster computers, the development of what I would call the super work station, and national networks. I think faster computers will continue to result from faster single processors.

In turn, I think those will be derived from either improving the current silicon technology, or going with new materials such as gallium arsenide, or perhaps the new superconducting materials and Josephson-Junction computers, or perhaps even more futuristic op-

tical or some combination of electrooptical computing. Those clearly offer significant increases in speed to current computers.

The Japanese claim to be developing, currently, the fastest single processors, and I think they are well-poised to dominate this area in the future. In the past few years, however, it has become increasingly apparent that new algorithms and massively parallel computer architecture offer the greatest potential for significant increases in speed, and I think the United States still leads in that area.

At Los Alamos we have concentrated much of our recent effort precisely in that direction, and that is, the developed distributed computing environment of supercomputers and work stations that are all hooked together with very high-speed networks, and networks that you essentially make transparent to the user.

So we have made much progress that has allowed us to solve complex problems, and at Los Alamos that means problems such as nuclear weapons safety, enhanced oil recovery, high energy physics, or even global climate change, and we have done this by developing new software, new algorithms, for the massively parallel machines such as the Connection machine, and again, the workhorses of most of what we do are the front-line Cray machines.

So let me turn now specifically to Senate Bill 1976, and as I mentioned, I think its greatest impact can be to keep the U.S. at the leading edge of computing, and what I would like to recommend is even more aggressive pursuit of the type of role that I just have explained, that Los Alamos has played, by developing several centers that I will call collaborative R&D centers for high-performance computing. The centers would promote collaboration between the Nation's leading computational centers and U.S. industry.

At Los Alamos we have in the past accomplished this type of role essentially entirely with the support of the nuclear weapons research, development and testing program. I think that is no longer possible today, in fact is no longer appropriate today, since our contributions covered the broad spectrum from let us say sequence mapping and sequencing of the human genome to enhanced oil recovery. So therefore, the support base must also be broadened.

Such collaborative R&D centers, as at least I envision them, would do the following things: first of all, they would continue to challenge computer manufacturers with demanding applications. They would identify requirements for suppliers and inventors. They would test newly-developed computational environments, including new architecture software, visualization techniques as well as networks, and they would also help to develop hardware, software and algorithms in conjunction and collaboration with computer vendors, and of course then they would help to disseminate that knowledge and experience gained.

By collaborative, I mean that in at least three sense of the word. That is, we would collaborate with universities to make certain that we would incorporate the latest ideas in high-performance computing. Secondly, we would collaborate with industry to set common goals and objectives. Then, third, by collaboration I also mean joint sponsorship between the government and industry.

A strong and deliberate collaboration such as the collaborative R&D centers is absolutely necessary, in my opinion, for us to stay

competitive. To rely on dated concepts of technology transfer spin-offs in my opinion simply will not work.

I think, as you well know, technology transfer is a body-contact sport. You have to collaborate on technology as it is being developed, rather than just handing off the baton in the rely race, so you have to do the collaboration.

We have developed a successful model for working with U.S. industry with the DOE superconductivity pilot centers that exist at Los Alamos, Argonne and Oak Ridge National Laboratories. These were established in 1988, and I think they could serve as a good model for the collaborative R&D centers that I mention here.

For these centers to be successful, they should be developed at laboratories that have a demonstrated track record of working with the computer industry. They must not only have state-of-the-art computing capabilities but also demonstrated need for large-scale computing power to solve important problems of national interest. I think Federal sponsorship of the collaborative R&D centers will provide U.S. industry the leverage that it needs to fund its own involvement in these endeavors.

So let me conclude by saying that as you have heard several times this afternoon, continued advances in high-performance computing are critical to national security, I think to our technological, industrial, as well as our economic welfare. Senate Bill 1976 provides an opportunity for the DOE and its labs to keep the United States at the leading edge of high-performance computing, and I suggest that the U.S. Government strongly encourage this closer relationship between the leading computing centers and users with the suppliers in industry.

So recognizing that the old approaches to technology transfer are really outmoded in this fast-paced global race, I encourage you to consider innovative mechanisms such as collaborative R&D centers to have us work with you, as industry, and with universities.

So let me also then conclude with a final comment that I think the high-performance computing program is most appropriate for government sponsorship, because after all, government is one of its most important customers and society at large is the most important beneficiary.

I thank you for the opportunity to speak, especially out of order.

[The prepared statement of Dr. Hecker follows:]

Comments on The Department of Energy High-Performance Computing Act of 1989 (S. 1976)

Statement to the Energy Research and Development Subcommittee of the Senate Energy and Natural Resources Committee

Siegfried S. Hecker
Director
Los Alamos National Laboratory

March 6, 1990

Mr. Chairman and members of the Senate Energy Research and Development Subcommittee, I am honored to be asked to appear before you to discuss Senate Bill S. 1976, "The Department of Energy High-Performance Computing Act of 1989." I am Sig Hecker, Director of Los Alamos National Laboratory.

EXECUTIVE SUMMARY

Senate Bill S. 1976 and the Bromley High-Performance Computing Program[1] stress the importance of high-performance computing to U.S. science and engineering as well as to U.S. society and productivity. The bottom line is that a country that seeks to control its future must lead in the development and application of high-performance computing.

The United States still leads the world in high-performance computing, but a national initiative is timely and necessary if we are to maintain that lead. Many problems of technical and societal importance are far too complex to be addressed adequately even with the advanced computational capabilities now available to scientists and engineers. But with concentrated and focused efforts, we will be able to continue the remarkable technological developments that started this revolution we call the computing age. We are currently at the doorstep of dramatic technological improvements that promise to change the way we do science and business, in fact the way we go about our everyday lives; but there is intense international competition for control of the capabilities key to this revolution.

Because these are times of tight federal budgets, the nation must leverage the existing federal R&D funding to encourage U.S. industry to continue and intensify development of high-performance computing technologies. We must collaborate with industry to ensure that their products meet the needs of the scientific community, and we must leverage federal R&D to support industry.

We foresee that supercomputer suppliers will continue to develop more and more powerful machines, extending to the trillion-operation-per-second range (compared with the billion range of today) and beyond. These machines will be critical in addressing not only the so-called Grand Challenge problems (as defined in the Bromley report[1]) but other important technical problems as well. High-speed networks and user-friendly interfaces will extend this powerful computing technology into offices and laboratories, and ultimately these capabilities will enter the mass market as the brains of a powerful multimedia center available to the home.

S. 1976 can help achieve this vision by making a national commitment to (1) maintain the leading edge of computing technology and (2) to make user-friendly supercomputer power available to a large segment of the U.S. population.

I want to emphasize three messages in my presentation today:

- *High-performance computing requirements should be driven by applications*—this will keep a focus on useful computing performance versus computing for computing's sake.
- *The Department of Energy (DOE) and its laboratories should play a leading role in this national initiative because they have important applications that require increased performance, and they have demonstrated the ability to run large computational centers and enhance useful performance.*
- *To encourage even greater interaction between the users and suppliers, I propose developing Collaborative R&D Centers (CRDCs) designed to keep the United States at the leading edge of computer performance through innovative collaborations among industry, the federal laboratories, and universities.* The Centers would develop joint goals and be jointly sponsored by industry and the federal government.

The Los Alamos National Laboratory, currently home of the most powerful scientific computing center in the world, stands ready to do its part toward these goals.

INTRODUCTION

We at Los Alamos support the High-Performance Computing Program by D. Allan Bromley and the principles of S. 1976, "DOE High-Performance Computing Act of 1989." High-performance computing is critical to U.S. national security and competitiveness. It has taken on a new and broader importance in recent years with extraordinary potential for future gain.

This initiative is timely because, although we are still the world leaders in high-performance computing, the environment surrounding supercomputing is changing rapidly as high-performance workstations and networks evolve as integral components of the computing environment. In addition, our lead continues to be eroded by Japanese advances in computing technology and systems and in manufacturing. As a nation we must respond with a new strategy that addresses these challenges. The High-Performance Computing Program is a significant step in this direction.

We support the intent expressed in S. 1976 that the DOE take a leading role in this initiative. The mission of the DOE has always required the application of leading-edge high-performance computing: first, to problems of national security and now to challenges of energy, the environment, and health. In pursuing its mission, the DOE has established in its national laboratories the world's most powerful scientific computational capability.

As a consequence, the DOE and its laboratories are in a position to help the U.S. maintain its leadership, foster the strengthening of the U.S. computer industry, and help disseminate high-performance computing into analysis, design, and manufacturing for U.S. industry. The laboratories are perhaps the most sophisticated, demanding users that push the requirements of high-performance computing to its limit; they have an unmatched record of close working relationships with computer vendors; and they have excelled in the development of high-performance computing technologies and computational science.

At Los Alamos, we have dedicated ourselves to these goals for over four decades. We, and other DOE laboratories, have the environment to lead a national effort aimed at enhancing the development of high-performance computing technologies in cooperation with U.S. industry and universities.

Before commenting specifically on the High-Performance Computing Program and the "DOE High-Performance Computing Act of 1989," I would like to explain the role Los Alamos has played in high-performance computing, what we are doing now, and how we see the future developing.

LOS ALAMOS AND SUPERCOMPUTING: A CASE HISTORY OF ACHIEVEMENT

Supercomputing originated with the "Los Alamos problem," as predicting the performance of the first atomic bomb was called during the Manhattan Project. The first attempts at supercomputing (a term relative to time) were made with arrays of mechanical desktop calculators organized by teams of operators. In 1945, Los Alamos researchers were the first demanding users of the University of Pennsylvania ENIAC, the world's first large-scale electronic computer.

Shortly thereafter, Los Alamos implemented the John von Neumann "stored program" architecture into its MANIAC computer, with which it ran the first large-scale hydrodynamic simulation of a thermonuclear weapon. The MANIAC was designed and built by Los Alamos scientists.

In 1953 IBM produced its first von Neumann machine with capabilities similar to those of the MANIAC, the IBM model 701. Serial number one was installed at Los Alamos. This began the era of Los Alamos collaboration with computer vendors that still continues today.

The first IBM connection lasted into the mid 1960s, at which time Los Alamos began to work closely with Control Data Corporation to improve its machines, which, in short order, became the fastest in the world. In time, Los Alamos had eight of Control Data's supercomputers.

The first Cray-1 was taken to Los Alamos by Seymour Cray for testing and evaluation in 1976 when it was the only asset of the start-up company, Cray Research, Inc. The close collaboration between Cray, the entrepreneurial supplier, and Los Alamos, the sophisticated and demanding user, was key to the success of Cray supercomputers and to the viability of the young company.

Since then, a succession of increasingly capable Cray supercomputers has provided the computational power needed by the Laboratory in ever greater quantity. The Cray-1 was followed by the Cray X-MPs and, most recently, by the Cray Y-MPs. The Cray Y-MP represents a twenty-to thirty-fold performance improvement over the Cray-1. From the first Cray-1 to the five Cray computers we have now (including two Cray Y-MPs), we have the equivalent power of about sixty-five of the original Cray-1 computers. In addition, Los Alamos is currently introducing the Thinking Machine's CM-2 into the production environment; presently there are two CM-2s on site.

Today, Los Alamos is the most powerful scientific computing center in the world serving more than 8,000 users throughout the nation via a national computing network. But it is our own computationally intense applications, starting with the complexity of designing nuclear weapons and now including a wide range of programs critical to emerging national and global challenges such as energy, ecological, and economic concerns, that require us to remain at the leading edge of computing technology. It has now become a commonly stated observation that computational science has taken its place with theory and experimentation as a third methodology for the conduct of scientific investigation.

The technology supporting supercomputing has grown rapidly. Over the past forty years, some applications have enjoyed an increase in computing speed of ten trillion, or ten orders of magnitude. This increase has come in almost equal parts from advances in computer hardware and from more sophisticated numerical algorithms. In the past few years, it has become increasingly apparent that new algorithms and massively parallel computer architectures offer the greatest potential for dramatic performance improvements. In addition, with increased speed has come the recognition of other significant requirements for large-scale computing:

- Real-time visualization,
- High-speed networking to provide convenient and powerful computing capability at your desk,
- Software to facilitate writing sophisticated scientific computer programs.

So, Los Alamos has turned much of its attention and collaborations in these directions. The high-performance computing model we are developing consists of powerful workstations at an individual's workplace connected to high-bandwidth and functionally transparent networks that provide connectivity to supercomputers. Using this scheme, we

seek to maximize the productivity of the individual using the integrated system. Hence, the task is one of integrating these different functionalities—workstations on desktops, networks, supercomputers—into a seamless, distributed environment.

We have made progress toward the integration of these technologies. As a major step, we have established an Advanced Computing Laboratory at Los Alamos. It was designed to be a catalyst for leading-edge research in these three key technologies. One goal of the Advanced Computing Laboratory is to help develop an integrated computing environment from these separate pieces. A current focus of this laboratory is the second-generation Connection Machine. It has 65,536 processors, 512 megabytes of memory, 4 frame buffers and monitors, and a 10-gigabyte data storage system. Let me add that the development of the Connection Machine was sponsored by the Defense Advanced Research Projects Agency (DARPA).

The Advanced Computing Laboratory also serves to foster interactions with the national computational science community and is the primary interface for Laboratory participation in several team activities: the DOE/Applied Mathematical Sciences Advanced Computing program, the NSF Science and Technology Center for Research on Parallel Computation (a consortium of Rice University, Caltech, Los Alamos, and Argonne), and the Gigabit National Test Bed. Los Alamos's Advanced Computing Laboratory is an excellent example of an existing center of excellence in high-performance computing that is positioned to participate in the development of computing technology through the High-Performance Computing Program. The activities of the Advanced Computing Laboratory are driven by the need to solve a broad spectrum of very computationally intense applications. I believe that it is a prototype for the Collaborative Research and Development Centers for high-performance computing, about which I will speak in a few minutes.

Other examples of Los Alamos's contributions to enhancing the new high-performance computing environment include the following:

- In 1986, a revolutionary algorithm for simulating fluid flows was proposed jointly by a Los Alamos and French collaboration. Classical models of fluid flow—partial differential equations—cannot be directly simulated on super-computers, but rather must first be discretized before the numerical simulation can begin. The new method, lattice gas hydrodynamics, is designed from first

principles to employ discrete processes and thus can be directly simulated on high-performance computers. Therefore, the lattice gas methods do not incur any of the traditional truncation errors associated with the classical approach and they can be computed at nearly optimal rates. These methods are particularly fast on the Connection Machine, because this hardware is very well suited to localized, bit-oriented computation. In addition to simulation of the Navier-Stokes equations, analogous techniques have been developed for flow in porous media, chemical reactions and combustion dynamics, plasma physics, and multiphase flow. This methodology, accompanied by modern computer technology, is allowing us to contribute to enhanced oil recovery.

- The goal of the Ultra High-Speed Graphics Project is to free scientists from the limitations of their current computing environment so they may more easily gain insight into the science of their problems and analyze the data from supercomputer calculations, which may run fifteen hours or more on a Cray supercomputer and produce vast amounts of data. These data are then reduced to a series of images that represent various physical variables calculated in the simulation. In the past it could take over fifteen minutes to display a single image. This project allows a scientist to view these images as a "real-time" animation. By using these techniques, scientists have recognized new physical features in older calculations, and subtle numerical errors have been removed. Without this capability, these effects would never have attracted attention and would have remained unnoticed. Indeed, scientists can interrogate a computer simulation by using visualization techniques in ways that are forever closed to experimentalists.

- Neural nets allow computers to learn tasks or patterns in the same way that humans learn. There are a number of problems—pattern recognition, for example—upon which traditional computer hardware and algorithms do not nearly approach human capabilities. Neural nets are formal, schematic models of these biological systems and have enjoyed success for certain applications. At Los Alamos, our current thrusts are to (1) understand the relationship between neural nets and other algorithms and (2) examine the use of neural net methods on real problems. To this end, we have written a compiler that can generate neural net application code on serial, vector, and parallel machines. On the CM-2, this compiler has generated some applications that have run at 3.2 gigaflops. Neural net

algorithms often perform an order of magnitude better than classical methods. One example of such an application is predicting transition to turbulence. We are developing a capability to employ these algorithms in gene sequencing, meteorology, astronomy, and seismology.

- The High-Performance Parallel Interface (HPPI) was conceived and designed at Los Alamos to satisfy our needs for very high-speed networking, including communications between supercomputers and real-time visualization of the results of scientific computations. The High-Performance Parallel Interface is a high-speed channel that will connect a variety of computing manufacturers' equipment into a computing network. Operating at 800 Mbit/sec (extendable up to 1600 Mbit/sec), the High-Performance Parallel Interface will increase the communication speeds between pieces of computer hardware by more than an order of magnitude over existing interfaces. This innovation appears destined to be an industry standard, with more than ninety computing vendors showing interest; some (such as IBM) are already marketing products based on the High-Performance Parallel Interface.

- Cray Research, Inc. (CRI) Joint Study. Beginning with serial number 1 in 1976, the goal of this collaboration has been to help continue U.S. preeminence in supercomputing and to provide unparalleled computational resources to Los Alamos scientists and engineers. On this first machine, incorporation of "vector recursion" as an announced feature and error checking in the Cray-1 memory were direct results of Los Alamos-Cray Research, Inc., interaction. Later, on the recommendation of Los Alamos, hardware scatter-gather capability was incorporated into the X-MP architecture to facilitate indirect array addressing. The Los Alamos HPPI standards were in part derived using the Cray HSX channel as a model. In the area of software, the CRI runtime scientific library incorporates ideas and some actual code originated at Los Alamos. This interaction has been very fruitful over the past fifteen years, with both parties benefiting tremendously. This collaboration is ongoing today with discussions centering on many of the issues of the High-Performance Computing Program.

- IBM Joint Study. IBM has shown a strong interest in reentering the competition for the high-performance end of the computer market, which is being threatened by Japan's NEC, Fujitsu, and Hitachi. The goal of the project from the Los Alamos

view is to involve another strong U.S. company in the high-performance computing marketplace. The IBM contribution to the joint study was its vast resources and knowledge in computer hardware and software development; the Los Alamos contribution was its variety of Grand Challenge problems and knowledge of the high-performance computing environment. The joint study has been very satisfactory to both parties. The IBM directions in scientific computing have been affected by their interactions with Los Alamos in several key areas, including compilers, floating-point representation, and communications. A major current thrust is in the high-performance developments in support of visualization. IBM recently announced the Los Alamos High-Performance Parallel Interface as a product on its 3090 series of computers.

- Thinking Machines Corporation Joint Study. In 1989, Los Alamos joined in a collaborative research effort with Thinking Machines Corporation with the 65,536-processor CM-2 in the Advanced Computing Laboratory. The focus of this collaborative research effort is the development of large-scale simulations on the CM-2 to drive understanding of this important computational architecture. Aside from the scientific results that have been achieved using this machine, we believe two changes in the basic philosophy at Thinking Machines are at least in part attributable to this interaction. First, the basic philosophy of computing on the CM-2 is currently under scrutiny because of the successes achieved using an alternate model developed by computational scientists at Los Alamos and other sites. Second, there is now a concentration on development of a scientific programming language capability (i.e, CMFortran) on the CM-2.

- The Gigabit Project, to be funded through the Corporation for National Research Initiatives, will demonstrate supercomputing using wide-area, high-speed (800 Mbit/sec) networks. The Los Alamos component of the Test Bed will be a network linking Los Alamos, Caltech, the Jet Propulsion Laboratory, and the San Diego Supercomputer Center. This project will investigate how to use multiple supercomputers and high-bandwidth, wide-area networks to solve important problems. Two research issues that arise are how to devise algorithms that hide delay inherent in long-distance networks and how to write or modify applications software to run in a manner distributed over computers separated by over 1000 miles.

THE FUTURE COMPUTING ENVIRONMENT

High-performance computing in the United States should be driven by two primary objectives: (1) we must stay at the leading edge of computing to solve increasingly demanding problems of science and technology, and (2) we must provide supercomputing power to all segments of society to provide new capabilities, products, or conveniences to business and the home.

Staying at the Leading Edge

In spite of the enormous improvements in performance noted already, the computational needs of the scientific community today are not being met. Indeed, every technological advance only whets the appetite and focuses the need for greater capability. Further, there are Grand Challenges of potentially enormous social benefit facing science that require a dramatic increase in the capability of high-performance computing.

Particle physicists yearn for immense increases in computational power to explore lattice calculations in search of new insights into quantum chromodynamics (QCD). The impact on the discipline could be profound, guiding detector development and experimental priorities at the Superconducting Super Collider (SSC). Atmospheric modeling requires vast improvements in spatial and temporal resolution for effective weather prediction. The international interest in sequencing the chemical nucleotides composing the genetic material of the human genome has increased dramatically. Such an endeavor would require vast computational capabilities in data management and analysis. Scientists hope to eventually understand condensed matter on an atomic level, which will allow the design of bulk materials with pre-selected properties. To be successful, this undertaking will require enormous improvements in supercomputer capability. Other examples can be found from virtually every discipline of science and engineering.

These problems require that we stay at the leading edge of computing technology. At Los Alamos, we depend on the fastest production supercomputers available in the world. For the last decade, these machines have been made by Cray Research, Inc. Cray Research has increased the performance of its computer systems by 20 to 30 times and has increased the memory size by over a factor of 100. These computer systems continue to be the workhorses of the Los Alamos capability. The recent integration of Thinking Machine's

CM-2 into the Los Alamos computing environment in a useful manner has already allowed us to make significant progress in several important areas of science and technology.

The simulation of the physical processes taking place in a nuclear explosion has historically been the bellwether technical problem driving the development of supercomputing, and it remains an extraordinary challenge. To simulate nuclear weapons behavior in three dimensions and with the desired phase-space resolution, orders-of-magnitude increases in all aspects of supercomputing capability are required. Simulations of nuclear explosions have become more important as we have reduced our reliance on underground nuclear testing and decreased the availability of experimental data. Hence, understanding the fundamental behavior of processes and materials under extreme conditions is more important. We have run our most sophisticated three-dimensional computer codes on the CM-2 and have gained new insight into nuclear weapons safety. The CM-2's parallel architecture and large memory allowed the problem to be modeled in a realistic three-dimensional manner.

In collaboration with Caltech and Thinking Machines Corporation, we have developed a *Lisp version of a state-of-the-art lattice QCD code for the Connection Machine. The code uses the hybrid Monte Carlo update method and employs the most efficient known method for inverting what is called the pseudo-fermion matrix, pre-conditioned, over-relaxed minimal residual. Physics results were reported from this code at the International Lattice Gauge Theory Conference in Capri, Italy. This analysis indicates the calculated mass of the rho meson, as a function of quark mass, may be approaching the physical regime for our current simulations, a first for this type of approach. A new version of this simulation is being developed in close collaboration with Thinking Machines; the new code will be implemented primarily in CMFortran rather than *Lisp, and we expect that the assembly language kernel of the code will execute at roughly 8 gigaflops.

We have also implemented a three-dimensional, explicit, free Lagrangian hydrodynamics code on the CM-2. This code is 99% parallel versus 95% parallel on a Cray Y-MP8/32. This improvement was realized because of hardware improvements designed into the CM-2 to support a parallelized indirect accumulation step that occurs often in our numerical physics algorithm. Recently, in support of global climate studies, Thinking Machines Corporation, the National Center for Atmospheric Research, and Los Alamos have collaborated to write a finite difference version of an NCAR fluid dynamics code that

runs at 6.1 gigaflops versus approximately 1 gigaflop on a Cray Y-MP8/32 for this one specific problem.

The projects listed here help us maintain our national security and promise social benefit, including energy-efficient automobiles, new pharmaceutical products and vaccines, and understanding global climate change. But for the promise to be realized significant advances in computational capability are needed by the scientific and engineering user communities. Presently intractable problems must enter the regime of affordable computer calculations. Computational problems that take major production time must enter the regime of routine calculations. Major improvements are needed in speed, memory, distributed processing, network bandwidth, visualization, user interface, and other related attributes and technologies. As quickly as practical, new experimental approaches to computer architecture must enter the production mode.

Workstations and networks are critical parts of the high-performance computing environment. Current supercomputers run at the rate of billions of arithmetic operations per second while national networks operate at maximum rates of 1.5 million bits per second or about 25,000 arithmetic operations per second. This mismatch of speeds precludes effective distribution of work among geographically distributed supercomputers and supercomputer users. Speeds approaching 800 million bits (800 Mbits) per second are required for effective distributed computing.

Pushing forward the state of the art in high-performance computing is not, therefore, a mere technological wish item of the scientific community; rather, it is a vital key to a better future for humanity.

Supercomputing Power for Everyone

Besides these leading-edge requirements, high-performance computing has yet to make its greatest impact on business and on the everyday lives of all of us.

Supercomputers and the associated technologies have been developed for scientists who are traditionally willing (in fact, eager) to use the most powerful and appropriate tool that will allow them to solve their problems. This approach has resulted in a collection of capabilities that are powerful for solving problems but are not always considered to be easy to use by a less determined audience. If high-performance computing capabilities

and concepts are to be integrated into the real world of business and leisure, a considerable amount of research and development will be needed to make these systems user friendly, so that the power of a supercomputer is accessible through the touch of a mouse on a desktop, either in the office or at home.

A computing environment that is friendly and available to large segments of our population would likely change the entire fabric of American life. Business decisions and operations already depend largely on computing, but complete connectivity and access to national data bases with powerful computers would increase their productivity and hence the nation's international competitiveness. Access to data bases could influence daily decisions that vary from health care and diagnosis to automobile repair. Already one automobile manufacturer is developing a capability that will centrally track all the problems that each automobile has ever had and help to diagnose them based on the complete history of other similar automobiles. That kind of capability should become commonplace in the future.

Powerful home computing capability could entirely change the way we educate our children and ourselves and enhance the effectiveness of the education. Imagine learning new languages or technologies or world events using a full-time, completely responsive, intelligent electronic tutor that resides in your home.

In a recent publication, Keyworth and Abell[2] described their vision of empowering the people. It provides for high-performance computing capabilities with access to high-speed, wide-area networks that will interconnect all parts of the United States like the telephone system does today. The family television, equipped with high-definition output and with the processing and memory capabilities of current supercomputers, would be the brain center. User-friendly operating systems and software will allow us to rapidly draw upon information generated or stored around the globe.

Such a vision may appear quite optimistic, but it is not inconsistent with Dr. Bromley's peek into the future published with Paul Huray.[3] Making available more and more computing power to more and more people with greater convenience and tied in a transparent fashion to the rest of the nation's scientific institutions, businesses, and homes will clearly enhance our productivity and standard of living.

THE HIGH-PERFORMANCE COMPUTING PROGRAM AND S. 1976

The vision of the future we have just outlined is consistent with the intent of the High-Performance Computing Program and S. 1976.

As summarized in the Bromley high-performance computing report and illustrated in the attached Figure 1 from the report, to enable the competitiveness of the United States in high-performance computing technology we need to develop a foundation using basic research and human resource development. On that critical foundation we need programs to (1) develop high-performance computing systems and advanced system software, (2) increase the resource-sharing capabilities across the nation by developing a National Research and Education Network (NREN), and (3) support Grand-Challenge research work through the development of software components and tools and algorithm developments.

I support a strong role for the U.S. government in this initiative to keep the nation at the leading edge of computing. Many of the scientific Grand Challenges and many of the most demanding technical problems are either directly a government responsibility or are of intense interest to society. The key technology problems include those in defense, space, environment, health and human welfare, etc.

S. 1976 properly points out that the DOE and its laboratories should play a leading role in maintaining U.S. preeminence in this critical area. Similarly other federal agencies such as DARPA, NSA, NSF, NASA, and others have and will continue to play a vital role. The larger government role is discussed in Senator Gore's bill S. 1067, "National High-Performance Computer Technology Act of 1989."

The federal government can also play a supportive role in bringing supercomputer power to the entire society. One of the most important functions will be to help provide the information infrastructure that will enable effective networking and distributive computing. The NREN proposed in the High-Performance Computing Program and S. 1976 is an important step in that direction. The government should encourage the rapid transition of ideas and technologies developed at the leading edge to flow to the mass user sector. If we stay on the leading edge and push forefront technology, then private industry will be the principal player in developing the futuristic "supercomputing power to everyone" environment that I described above.

The greatest contribution that Los Alamos (and several other DOE labs) can make to the High-Performance Computing Program is to shape the leading edge of computing. In the past, we accomplished this role almost entirely with the financial support of the nuclear weapons research, development and testing program. That is no longer possible, nor appropriate, today. As mentioned above, our contributions now cover a broad spectrum from the human genome project to enhanced oil recovery. From this application base we wish to enhance the most dominant trends that we see in computing today: (1) the continued push towards faster computers, (2) the evolution towards the super workstation system, and (3) the development of national network technology. All will require the development of advanced software and algorithms to turn these trends into useful advances in computing performance.

To substantiate the first of these trends, I would remind you that in just a few short years, the performance of supercomputers increased from the million floating-point operations per second, or megaflop range, to the billion, or gigaflop range. We can now confidently project machines capable of performing a trillion or more arithmetical calculations every second, so called teraflop machines, within the next five years.

Leading the way toward such incredibly fast machines are advances in the materials and other means of transmitting information. Here, I would point to the possibility of superconducting Josephson-Junctions. These will significantly increase speed and lower the cooling requirements of supercomputers resulting in reduced size, cost, and weight. The incorporation of gallium-arsenide circuitry as planned by some supercomputer manufacturers also holds out the promise of greatly improved performance because gallium arsenide boasts not only lower power consumption and higher temperatures than silicon, but more importantly it promises a switching speed three times as fast as that of silicon.

In addition to squeezing more and more speed out of the processors and their associated circuitry, we have already mentioned that a simultaneous revolution is occurring in the approach to information manipulation itself, one that depends upon an increased number of processors to process data in parallel rather than serial form. Massively parallel machines, such as the CM-2, are making great strides, and we have demonstrated that for certain problems they have significant advantages in speed.

The evolution of personal workstations with the capabilities of today's most advanced supercomputers is increasingly likely because of advances in architectures and fabrication techniques. These advances may provide feature sizes on the order of one-tenth of a micrometer in integrated circuits—a trend leading to a product that many refer to as a "Cray on a chip." These workstations will provide a superior interface with other computing resources and will become a multimedia work center, incorporating tools that scientists need for all aspects of their work.

COLLABORATIVE R&D CENTERS FOR HIGH-PERFORMANCE COMPUTING

How can S. 1976 best promote these important trends and keep the United States at the leading edge? The DOE labs and universities can help develop networking technologies, new algorithms, software, and hardware in quest of their scientific challenges or technical applications. The current bill provides support for these activities.

Based on our vision of the future computing environment and our past experience, I believe that the bill should explicitly encourage stronger collaboration between this nation's leading computational centers and U.S. industry. I suggest that substantial resources be devoted to establishing a small number of Collaborative R&D Centers for high-performance computing aimed at bringing together the sophisticated users (federal labs and universities) and the suppliers (U.S. industry).

A strong, deliberate collaboration is necessary to remain competitive. Relying on dated concepts of technology transfer spin-offs will not work. We have learned that effective technology transfer is a body-contact sport (collaborating on technology as it is being developed) and not just handing off a baton in a relay race. The Japanese can now build single processors as fast as anything we have in the United States. The next few years definitely favor even more rapid Japanese advances because of their dominance in the memory chip market and new materials such as gallium arsenide. The United States depends on better integration of processors into systems or networks to preserve our lead in high-performance computing. The principal objective of the Collaborative R&D Centers would be to enhance this advantage by encouraging even closer collaboration between users and suppliers.

In addition, the Collaborative R&D Centers could also address collaboration in new technologies that may allow us to leapfrog the current Japanese technological

28-910 O - 90 - 5

advantages. These technologies include advanced photolithography using XUV-free-electron lasers (being considered at Los Alamos), optoelectronic computing, and computers built with superconducting Josephson-Junction components. Programs at DOE labs and universities in these areas would offer substantial leverage for U.S. industry in a collaborative R&D effort.

To meet the objectives of the High-Performance Computing Program and S. 1976, the Collaborative R&D Centers, or CRDCs, would function similarly to the way Los Alamos and a few other DOE labs have worked in the past with the computer industry. The Collaborative R&D Centers would

- *Challenge computer performance and computing tools with demanding applications.* In the same way that the DOE laboratories have provided the "pull" for high-performance computing technology from the computing industry in the past, the CRDCs would challenge high-performance computing vendors to continue pushing the state of the art.

- *Identify requirements for suppliers and inventors.* Just as the DOE laboratories have shown high-performance computing vendors the computing environment that will exist in the future, CRDCs, with their demanding users and applications, would help high-performance computing vendors understand what products will be needed next.

- *Test this new computational environment, integrating new architectures, software, visualization.* The CRDCs must collaborate among themselves as a distributed center for testing the national computing environment. Several examples are the DOE Applied Mathematical Sciences effort, the NSF Science and Technology Center for Research on Parallel Computation, and the National Gigabit Test Bed.

- *Develop software, hardware, algorithms, etc., in collaboration with high-performance computing vendors.* Software and hardware development often benefit immensely from collaborations. Some software, for example, has proven difficult or impossible for vendors to develop because they typically do not have the demanding applications in house. Innovative hardware ideas often originate with users that are closest to the problems. Vendors are then invited to collaborate

in the developments, and the resulting products may become commercially viable. An example is the HPPI cross-bar switch originally conceived at Los Alamos, prototyped with Digital Equipment Corp., and now marketed by Network Systems Corp.

- *Develop and test new distributed networking systems.* The CRDCs will undoubtedly build high-performance computing environments consisting of hardware and software from many different vendors. These equipment and networking concepts will be tested on demanding applications.

- *Disseminate knowledge and experience gained to other users.*

Like the highly successful DOE Superconductivity Pilot Centers established at Los Alamos, Oak Ridge, and Argonne National Laboratories in 1988, the CRDCs should be collaborative centers in three senses of the word: (1) they should work with universities to incorporate the latest ideas into high-performance computing; (2) they should work with industry to set common goals and objectives; and (3) they should be jointly sponsored by the government and private industry.

The Pilot Center model has demonstrated that industry is enthusiastic about collaborations when it can leverage its investment and resources. For the CRDCs to be successful, they should be developed at laboratories that have a demonstrated track record of working with the computer industry. They must not only have state-of-the-art computing capabilities but also the demonstrated need for large-scale computing power to solve important problems of national interest.

Federal sponsorship of the laboratories and universities to perform the functions of the CRDCs outlined above and to solve their computationally intense problems is imperative to give industry the leverage it needs to fund its own involvement in the joint programs of the CRDCs. This mechanism of working with industry will also provide a greater assurance that technology developed in the United States will benefit U.S. industry rather than being brought to market by others before U.S. industry can respond.

CONCLUSION

In conclusion, continued advances in high-performance computing are critical to the national security and economic, industrial, and technological welfare of the United States. Other nations have recognized this fact and have formulated coherent strategic plans aimed at gaining dominance in this area. Their technological advances and national foresight and persistence are even now eroding the U.S. lead.

We can no longer afford to study the problem. The time has come for concerted action. Because the time is short and the funding not unlimited, we must act to strengthen our lead in high-performance computing by building on existing excellence. Let us avoid a costly learning curve and leverage our limited resources by using those with a track record of working on problems of overriding national and international importance. Recognizing the need for cooperation, let us encourage closer interactions between computer users and suppliers. Recognizing that the old approaches to technology transfer are outmoded in a fast-paced global race, let us give great consideration to innovative mechanisms—mechanisms like the Collaborative Research and Development Centers—that will facilitate collaboration among national laboratories, U.S. industry, and universities. Los Alamos stands prepared to play a central role in this important national effort.

REFERENCES

1. *The Federal High-Performance Computing Program*, Executive Office of the President, Office of Science and Technology Policy, September 8, 1989.

2. George A. Keyworth, II, and Bruce Abell, *Competitiveness and Telecommunications*, The Hudson Institute, Indianapolis, IN, 1989.

3. D. Allan Bromley and Paul G. Huray, "The U.S. Government's High-Performance Computing Program: How Much Will We Use It to Leverage America's Future?" *Roll Call*, January 29, 1990.

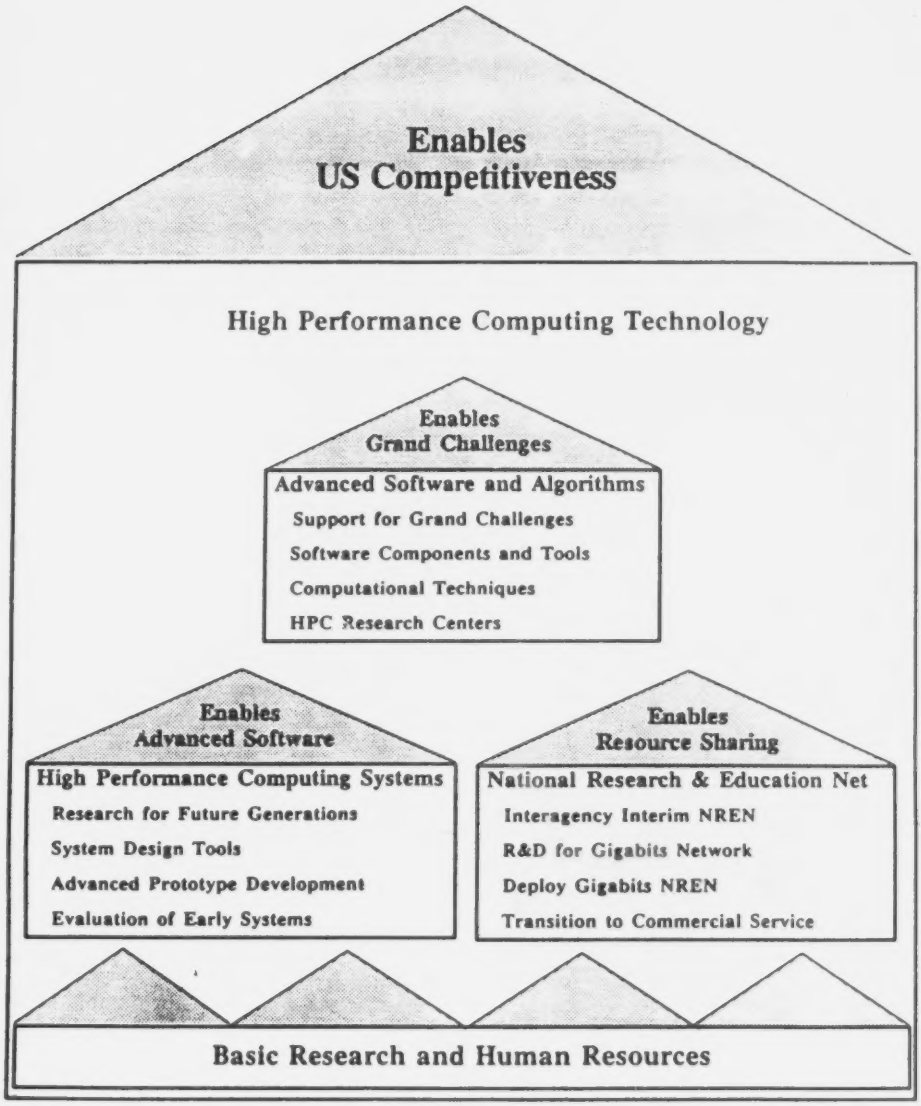

Fig. 1 – Relationship of HPC Program Components

(The Federal High-Performance Computing Program, Executive Office of the President, Office of Science and Technology Policy, September 8, 1989)

Senator FORD. Thank you, Dr. Hecker.
Dr. Connolly.

STATEMENT OF DR. JOHN W.D. CONNOLLY, DIRECTOR, CENTER ON COMPUTATIONAL SCIENCES, UNIVERSITY OF KENTUCKY

Dr. CONNOLLY. Good afternoon, Mr. Chairman.

Senator FORD. We always need a clean-up hitter.

Dr. CONNOLLY. I am very pleased to be here today to testify on behalf of the Department of Energy High Performance Act of 1989. I am the Director of the Center on Computational Sciences at the University of Kentucky, which is the center dedicated to the use of supercomputing as a tool for the solution of science and engineering problems.

Before joining the University of Kentucky, I was the Director of the Office of Advanced Scientific Computing at the National Science Foundation. This was the office responsible for the establishment of the NSF national supercomputing centers and the national network known as NSFnet.

In this previous position, I was the NSF representative on the Science & Engineering Computing Subcommittee of the FCCSET, and as such I was involved with initial phases of the study which led to the OSTP report which we have heard so much about today.

I have therefore been an advocate of a focused U.S. effort in supercomputing in networking for several years from both the Federal and the state perspective, and since I no longer work for the Federal Government I am allowed to say that I am very much for this bill today.

Senator FORD. Oh, we finally get to the truth.

Dr. CONNOLLY. I would like to make four points in the brief time that we have left here. One is the importance in this area of government-university-industry cooperation. Several of the witnesses have mentioned that.

I think it is amusing to note that exactly 100 years ago it was a government and university project that developed the first electric computer which was used to count the 1890 Census. That was a professor at MIT working with the Census Bureau, and the first electronic computer—it is a well-known story—was built as a collaborative university-industrial project.

In the years since, as the computing industry has exploded, government funding has made possible many of the most important advances in computing: time-sharing, interactive networking, vector supercomputing. The role of government has historically been to encourage innovation and experimentation which benefits the entire industry.

The second point I want to make is that it is important—and I think this bill is a good example of it—it is important to have a balanced plan. When you read about computers in the media you often see that the speed and capacity of computer hardware is the only thing they speak about. As it has been pointed out many times before, gigaflops are a necessary but not sufficient requirement for a high-performance computing plan. Without a large effort in software and algorithm development, and networking

access to the hardware, and an investment in education and training, the hardware advances will be wasted.

So the current U.S. leadership position is based on a combination of these factors, and especially on a continued source of talented specialists produced by American universities. The continued competitiveness of the U.S. computing industry depends on the continued support of government funding.

The third point I would like to make is the importance of a coordinated plan for Federal high-performance computing. My experience with the NSF initiative in computing and networking, which has been built up over the last five years, has been extremely successful in providing state-of-the-art access for the university community. In fact, much of the NSF program was modeled on the programs developed in Department of Energy laboratories.

The success of the NSF plan was due in no small part to the fact that it was a top-down plan. The plan was funded off the top of the NSF budget, rather than being spread out over the many different individual NSF programs. Distributed funding of computing always leads to distributed systems.

The success of the Internet, which grew out of NSFnet and which is rapidly becoming a global information resource, has also been due to a considerable amount of effort invested in coordination from the top.

Similarly, high-performance computing at the University of Kentucky has been from the beginning a top-down plan, involving a separate budget for a supercomputing facility and for the multi-disciplinary Center for Computational Sciences. In fact, I do not know of any successful supercomputing or networking program which was not set up in this manner. Therefore, I think an important part of the Federal program should be a mechanism for coordination from the top so that the funds are not wasted in duplicative and incompatible systems.

I do not know what the best form of this mechanism is. I know I have some concerns about some of the mechanisms suggested, especially the option of using a FCCSET committee exclusively for the management of a Federal high-performance computing program. It bothers me a little. I used to be a member of that committee and I know what it can do and what it cannot do, and I mention that further in my written testimony.

My fourth and last point is that it is extremely important, as other witnesses have noted, to stress the importance of computational science. The ultimate justification of high-performance computing is to solve the important problems of science and engineering, the so-called grand challenges. Even though a computer is a fascinating machine, one needs to remember that the computer is a tool, which is a means to an end.

There are several excellent examples listed in the recent OSTP report and in this bill. I will not repeat them all. I have mentioned a couple and expand on them in my written testimony. Let me just mention one—the environmental modeling.

To take a program at random, the Department of Energy atmospheric radiation measurements program, the so-called ARM program, will be focused on gathering data with a feedback loop into computer models. This program, when it gets up to speed, will be

generating hundreds of tera bits of data per day. A tera bit—that is a million million bits per day, and sharing these data will require a high-speed national network much faster than the one we have today.

This is one of the things that this bill does. It gives us a network within five years which is a thousand times faster than the one we have now. Processing it will require fast supercomputers, much faster than we have today. Incorporating the data into predictive modeling codes will require software development by an increasing number of computational atmospheric scientists. It is clear that the ARM program will not achieve its goals without a greatly enhanced high performance computing infrastructure in all of its four aspects.

I mention some other examples, like materials research—dear to my heart, and to Sieg's heart—the superconducting materials, semiconducting materials; these future industries will all depend on having a substantially enhanced high-performance computing infrastructure.

Another very, very interesting example, although it is nonindustrial and not practical, is high energy physics. As a matter of fact, a substantial part of the National Science Foundation supercomputer and the University of Kentucky supercomputer is focused on trying to understand fundamental interactions between physical particles at the very smallest level. The people in this area tend to be the computational pioneers who are out there on a limb ahead of everybody else.

There are many other examples. I will not go into them, just the fact that there are many, many more opportunities opening up. The opportunities are opening up much faster than the funding to support them. The future of these expanded computer opportunities will probably happen whether this bill is passed or not, but the high-performance computing program will significantly accelerate the process and ensure that it will happen here in the United States.

Thank you for this opportunity. I would be happy to answer any questions.

[The prepared statement of Dr. Connolly follows:]

TESTIMONY TO THE SENATE COMMITTEE ON ENERGY AND NATURAL RESOURCES

MARCH 6, 1990

JOHN W.D.CONNOLLY

UNIVERSITY OF KENTUCKY

I am very pleased to be here today to testify on the Department of Energy High Performance Computing Bill of 1989.

I am the Director of the Center on Computational Sciences at the University of Kentucky. This Center is dedicated to the use of supercomputing as a tool for the solution of science and engineering problems. The faculty and students in this Center are involved in a wide range of multidisciplinary problems, from the understanding the fundamental processes of matter and the universe to the practical problems of medicine, economics, energy and the environment. We are particularly interested in innovative uses of supercomputers, and the interdisciplinary transfer of computational techniques. The Center is supported through funds from Federal, State and industrial sources.

I am also a theoretical physicist, so my testimony will be biased in the direction of how the initiative described in this bill will help basic research in science and engineering.

Before joining the University of Kentucky, I was the Director of the Office of Advanced Scientific Computing at the National Science Foundation. This Office was responsible for the establishment of the NSF National Supercomputing Centers and the National Network known as NSFnet.

In this previous position, I was the NSF representative on the Science and Engineering Computing subcommittee of the Federal Coordinating Committee for Science Engineering and Technology (FCCSET). As such, I was involved with the initial phases of the study which lead to the 1989 OSTP Report on the "Federal High Performance Computing Program". As you know, this report recommended a program very similar to the one contained in this bill.

I have therefore been an advocate of a focused US effort in supercomputing and networks for several years, both from the Federal and State perspectives. I would like to commend this committee for its support in the past, and for its foresight in crafting a balanced plan in this bill. I would like to see it succeed and will now make few comments on its behalf.

1) **The importance of government/university/industry cooperation:**

Historically, many of the most important advances in computing have been achieved through cooperation between the Federal government, universities and industry. Exactly one hundred years ago, it was a joint government/university project which developed the first electric computer. Professor Hermann Hollerith of MIT was awarded a contract to build a machine to collate data for the 1890 census. He later discovered many other uses of his machine and established his own company, which eventually became the beginnings of the International Business Machines Corporation.

56 years later, in 1946, it was another government-university collaboration which built the first electronic computer--the ENIAC, designed by faculty at the University of Pennsylvania to calculate shell trajectories for the Aberdeen Proving Ground.

In the years since 1946, as the computing industry has exploded, government funding has made possible many of the most important advances, such as time-sharing, interactive networking and vector supercomputers.

The role of government in the computing industry has historically been to encourage innovation and experimentation which benefits the entire industry. (Industry seldom invests in basic research which will help their competitors as much as themselves. No one but IBM and some of the larger Japanese companies can afford it.)

2) **The importance of a balanced plan:**

Much of the recent attention in the media and elsewhere has been focussed on the tremendous gains which have been made in the speed and capacity of computer hardware. As has been pointed out many times before, gigaflops are a necessary but not sufficient requirement for a high performance computing plan. Without a large effort in software and algorithm development, in networking access to the hardware, and an investment in education and training, the hardware advances will be wasted. The current US leadership position is based on a combination of these factors, and especially on a continuous source of talented specialists produced by American universities. The continued competitiveness of the US computing industry depends on the continued support of governmental funding sources.

The Department of Energy, and many other Federal and State agencies, have done a great deal to support high performance computing. They have all expressed publicly that High Performance Computing is a high priority. However, one has to wonder why such a high priority program did not do better in the FY1991 budget request.

3) **The importance of a coordinated plan:**

The NSF initiative in supercomputing and networking, which has been built up over the past five years, has been extremely successful in providing state-of-the-art access for the university community. Much of the NSF program was modeled on the programs developed in DoE laboratories. In fact, one of the goals of the NSF program was to bring access for university researchers up to the level of the DoE laboratories. The NSF centers have done a magnificent job in doing this, and indeed in some areas, may have surpassed the DoE laboratories.

The success of the NSF plan was due in no small part to the fact that it was a "top-down" plan. The plan was funded off the top of the NSF budget rather than being spread out over the many different individual NSF programs. Distributed funding of computing always leads to distributed (and therefore incompatible) computer systems. The success of the internet, which grew out of NSFnet, and which is rapidly becoming a global information resource, has also been due to a considerable amount of effort invested in coordination from the top.

Similarly, high performance computing at the University of Kentucky has been from the beginning a "top-down" plan, involving a separate budget for a supercomputing facility and the multidisciplinary Center for Computational Sciences. The latter provides a coordinated focus for the development of algorithms, software, networking, training and scientific applications. If the budget were to be distributed among the various departments, the whole effort would fail.

In fact, I do not know of any successful supercomputing or networking program which was not set up in this manner.

Therefore, an important part of a Federal program in High Performance Computing should be a mechanism for coordination from the top, so that the funds are not wasted in duplicative and incompatible systems. I do not know what the best form of this mechanism should be. There are several choices, such as the FCCSET which has been successful in coordination of some areas, or an independent entity like a public corporation, or by designating one agency to be take the lead in the various aspects.

I am somewhat concerned about the option of using a FCCSET committee exclusively for the management of a Federal High Performance Computing Program. While this committee is very capable of making policy decisions and guiding initiatives through the budget and formulation process, its structure cannot perform the day-to-day practical direction and management needed to carry out such a program. It would not be able to respond as quickly and flexibly as needed. The High Performance Computing program, in breaking new ground, will certainly encounter problems and glitches which characterize such efforts. There should be an organizational structure which can ensure that there is daily

attention paid to the administration of the program. Perhaps some entity should be formed to support the activities of the FCCSET and ensure the integration of the program.

4) **The importance of computational science:**

The ultimate justification of high performance computing is to solve the important problems of science and engineering, sometimes known as the "grand challenges". Even though a computer is a fascinating machine, one needs to remember that the computer is a tool which is a means to an end.

There are several excellent examples of grand challenges listed in the recent OSTP report and in this bill. I will not repeat all of them here, but I would like highlight some of those which I feel are the most urgent.

Consider the case of environmental modeling, which is quite justifiably listed as the first of the grand challenges. In the near future, many important policy decisions will have to be made on the basis of our best predictions of the consequences of the pollutants we are putting into our environment. These predictions are all based on supercomputer modeling of a very complicated system. Because of the limitations on computing capacity, the models are necessarily approximate. Large scale effects are quite well handled, but entities (like clouds) which are smaller than a few hundred kilometers, are included only crudely.

Improvement of the predictive models will depend on a better understanding of many different disciplines,- physics, chemistry, fluid dynamics, ecology, hydrology, geology, biology and economics, to name a few. Better understanding will lead to better computer codes, but at a cost of greater complexity, which in turn will require better computer resources. In order to improve the resolution of a global model by a factor of 10, a factor of 10,000 is required in computational power. Some part of this will come from hardware speed-up, but a significant fraction will also come from algorithm and software improvement.

Programs like the DoE ARM (Atmospheric Radiation Measurements) program will be focussed on gathering data with a feedback loop into the computer models. This program when it gets up to speed will be generating 100's of Terabits of data per day. Sharing these data will require a high-speed national network, processing it will require fast supercomputers, incorporating the data into the predictive modeling codes will require software development by an increasing number of trained computational atmospheric scientists. It is clear that the ARM program will not achieve its goals without a greatly enhanced High Performance Computing infrastructure in all its aspects. In this field, it is just as important to analyze and disseminate the data as to gather it in the first place.

A second important area is materials research. The entire electronics industry depends on the understanding of the formation and performance of semiconductor materials. The current generation of supercomputers has got to the point where microscopic simulations of semiconductors are becoming useful for actual device design. At the "nanostructure" level, where the size of an electronic element is about one-millionth of an inch, the ordinary laws of electronics break down and the laws of quantum physics take over. At this level, only a supercomputer can give accurate predictions of the device behavior. The design cycle of future electronic industries will be strongly dependent on supercomputer simulation.

A new, and as yet embryonic, industry will soon develop around superconducting materials. The theory behind the new family of high temperature superconductors is also embryonic, so that the understanding of how these materials work is incomplete. But it is quite clear that this understanding will be formulated in terms of high performance computing simulations, and they will be just as important to the superconducting industry as they are for semiconductors.

Yet another industry will arise in biological materials, which are just beginning to be understood and catalogued through the use of high performance computing resources.

As a last example, I would like to cite a non-industrial, very fundamental application of supercomputers, namely, quantum chromodynamics, which is physics at the smallest level. The researchers in this area are the "computational pioneers",- the first to exploit such advances as vectorization, parallelism and special processors. This field has no known practical application as yet, but the results will be very important to the understanding of the data produced by the Superconducting Supercollider. Also, it attracts some of the best students who often migrate to other areas, including the computing industry.

There as many other examples, equally useful or fundamental, on which researchers at laboratories everywhere are concentrating. The computational field is attracting more and more scientists, because of opportunities like these examples, and because of the exponential increase in computer technology more and more opportunities will open up.

The future of greatly expanded computer opportunities will happen whether or not this bill is passed, but the establishment of a Federal High Performance Computing Program will significantly accelerate the process and ensure that it will happen here in the United States.

Senator FORD. Thank you very much, Dr. Connolly. I am going to start with Mr. Rollwagen and ask him a question or two and some things that have bothered me in trying to prepare myself for this.

I understand that both of your companies are dependent on Japanese for components for your computers. To what extent are you dependent on them?

Mr. ROLLWAGEN. That is a changing situation, senator. In fact, I am very pleased to be able to tell you that as of last year, for the first time in a decade we have been able to make our most advanced computers from all U.S. parts, so that in fact today, we can manufacture Cray systems with all U.S. parts.

We do not choose to do that exclusively because Japanese suppliers are important to us.

Senator FORD. So you do use Japanese suppliers?

Mr. ROLLWAGEN. Absolutely, we do. But it is very important to recognize that we actually fabricate, now, many of the parts internally in our own company and we have developed U.S. sources of supply in all areas, and particularly in semiconductors, so that we can buy all the parts in the United States.

Senator FORD. So you are saying to this senator that you are not dependent on Japan, or that you would not be dependent on it?

Mr. ROLLWAGEN. Not in the ultimate sense. We probably cannot get the entire volume that we need to support all of our production in the United States, but we also cannot be cut off at the pass, if I can put it that way.

Senator FORD. Any concern about the Japanese learning certain technology with the trade, with our American companies, or our relationships back and forth between the two entities?

Mr. ROLLWAGEN. Almost the reverse. I think that, clearly, our toughest competitors at the high end of this business are in Japan, three companies in particular, and obviously the more they learn about supercomputing and how to put these machines together the more formidable they become.

On the other hand, when I look at our own technical progress, to a great degree this is dependent on our expanding user base, and the fact that we do business in Japan probably benefits us more than the fact that our competitors are in Japan and therefore can have a closer look at our machines, let us say.

Because the industrial sector, in particular—in the automobile area, for example, the most advanced developments using our machines are being done in Japan at the Japanese automobile manufacturers, and as a result we have insight into what they are doing and this places demands on us similar to what Los Alamos continues to do as well, for improvements in our product, and we get ideas for how to move forward.

So in fact the international trade I think in this sense benefits us technically as well as economically, although I am not sure I am getting at exactly what may be behind your question.

Senator FORD. I am not trying to lead you down a primrose path, but I understand that Cray was surprised to learn that the Japanese company that was furnishing—that supplied semiconductors released a new supercomputer with key parts that Cray had never heard of before.

Mr. ROLLWAGEN. Yes, we have had——

Senator FORD. Well, let me just finish my thought here. How well can the Japanese delay our access to the state-of-the-art technology?

Mr. ROLLWAGEN. The Japanese in the semiconductor area have some unique capabilities that do not exist in the United States, and as a result they can control our access to that technology and we want to be sure to avoid that kind of situation. That is why we have worked hard to develop U.S. suppliers and multiple suppliers in every area of our business that we can.

We have had the experience in the past. I think it was a memory part was in use in a competitor's machine before we knew about that parts availability, so that was a little yellow flag, a caution flag for us in dealing with that particular supplier.

Since that particular event, we have not had a similar experience, but it is something that we have to be concerned about and we think the best defense is obviously to have alternative sources of supply.

Senator FORD. You are enthusiastic about your business, and it is growing and being competitive, and you find that the bureaucrats are trying to restrain growth. They are very slow and very careful. And while the rest of the world passes them by, they are still in a six-by-six with a green shade on, and a 60-watt bulb.

So let me ask you this: How significant is our lead in supercomputing today as compared to five or 10 years ago?

Mr. PALUCK. Well, I think the gap has certainly narrowed. I think right now the closest competitors, of course, are the Japanese. And I think they are within just a few years.

When you talk about supercomputing, you are talking about several different things; certainly hardware, which is just the raw speed of how quick can you make the machine go. And in that regard, they are in pretty good shape. They are catching up very, very quickly.

The other key area, though, is in the area of software. And there are two aspects to software, systems software and the area of compilers and operating system. And then the other is application software.

Here I think the U.S. has a much greater lead, except the thing that worries me is that the Japanese are excellent at being very pragmatic and plodding along, step by step by step, and, over time, it does not take very much time, where they continue to incrementally improve what they are doing, and all of a sudden you find out that they really have a superior product.

So I think they are very close. I think we are talking a year or two years. Which says that, and as we have learned, and I have certainly learned in the semiconductor industry, leadership is a very fleeting thing. And the only way to keep leadership is you have got to run faster than the guy who is trying to catch you.

And in that regard, we need all the help we can get. Not subsidies or anything like that, but just all the stimulation, all the encouragement that industry, but especially government, can provide. Because anything that slows you down just helps the other guy try to catch up.

Senator FORD. Well, what disturbs this Senator is I find that we are slipping further and further behind or the gap is getting closer

and closer in almost everything we do. And instead of being number one, we are beginning to look over our shoulder and see something, or we begin to look at those out in front of us.

Mr. PALUCK. Can I comment on that?

I do not think there is anything unique about supercomputers. I think that one of the things that is really changing right now that did not occur in the auto industry, it did not occur in the semiconductor industry, and in other areas where the Japanese did overcome American manufacturers, and that is, there is a very high awareness today, the fact that we are here talking about it today, but many other things. There is a very high awareness of the threat that exists.

The American entrepreneurism, the American enthusiasm, and innovativeness, in the right environment, can still succeed. And so, yes, there is a threat. Yes, we are being threatened. But there is reaction now. And that is different. That is what is different today than in other industries.

So it is not too late.

Senator FORD. I heard one of our senators make a speech the other day, and he talked about the difference between Japan and America. He talked about them going to school 280 days and we are going to school 180 days. They have homework and our kids come home to Nintendo, which they made. [Laughter.]

Senator FORD. So it does kind of grate just a little bit. And basketball, I watched it all weekend. [Laughter.]

Senator DOMENICI. That is all right, you are already smart.

Senator FORD. No, no, I am just waiting for you guys to figure it out and tell me.

You mentioned that buying practices by the Federal Government should be fair. Are they now or should they be changed?

Mr. PALUCK. Well, it gets back to what I just mentioned about we need all the encouragement we can get. And a retarding in any way is not good, because the other guy could catch us.

And so what comes to mind is really that there just needs to be a level playing field here in the United States. Just that there needs to be, and we certainly asked for it in Japan. And one aspect of the proposal of the bill is it could provide an open door for other companies into the Department of Energy, and that is outstanding and very, very healthy. And then it is up to those individual companies to prove themselves.

I think the practice of sole-source procurement is not a healthy thing. And I think all government procurement should be scrutinized and open to open bid, just like other things are done.

Senator FORD. We have a lot of sole-source contracts in the government.

Mr. PALUCK. There are situations.

Senator FORD. And that reduces competition, also, when you have sole-source.

Mr. PALUCK. Correct.

Senator FORD. Dr. Hecker, I am going to pass you, because you have got two colleagues here that are getting ready to ask you some questions for the record. And I think, rather than just probably asking the same thing, I will let them do that.

So I am going to slip past you then to Dr. Connolly. And I have only a question or two for him.

Doctor, how much access do universities have to the Department of Energy's supercomputers?

Dr. CONNOLLY. Very little. The Department of Energy computers are largely for Department of Energy purposes. Now, of course, they have such a large resource that, of the small percentage that does go to universities, it is very helpful.

There are particular areas of the Department of Energy supercomputer centers which are very much available to universities. And particularly in magnetic fusion and in some aspects of high energy physics.

But for the general university community it is very difficult to get into a Department of Energy laboratory.

Senator FORD. Doctor, I made a little note here, do we need to amend this bill to ensure education access fully? Do you feel it is necessary, so we will be sure that the university community and the research community and so forth would have access without having the door shut on them, or whatever?

Should we have some different type of language in this bill?

Dr. CONNOLLY. The Department of Energy does do some significant things for universities. But I would like to see more emphasis. I would like to see more emphasis in this bill on university access.

One of the important aspects of this bill is in the development of human resources. And in the computing field, human resources get developed in the university atmosphere.

Every one of the top flight scientists at DOE laboratories was at one time a university student. And if we are thinking ahead to the future, it is absolutely essential that we look at the raw material that is coming out of the undergraduate and graduate schools of this nation.

Senator FORD. Thank you very much, Doctor.

Senator Domenici.

Senator DOMENICI. Thank you very much, Mr. Chairman.

Let me just say to the four of you that, from this Senator's standpoint, this is the kind of exciting thing we ought to be doing around here, having these kinds of hearings and trying to accomplish some of the goals you have set out.

But I am kind of, on the one hand thrilled, and on the other hand a little bit worried.

First, let me throw something out, and this is particularly to the two industry people, and I am not so sure that either of you are typical of American industry, versus the national laboratories or the national government's scientists, because you are very intimately involved.

But, Dr. Hecker will remember that, starting about four years ago, when we started talking about technology transfer in the Department of Energy, and, in particular, from the three nuclear deterrent laboratories. And in a sense we are really talking about technology transfer here. The overall umbrella, as I see it, in the broadest sense of the word.

It is amazing, at the point we started, maybe four or five years ago, talking about technology transfer inhibitions within the Department of Energy, American industry was not at all interested at

that point in the national laboratories—and I am just using them today—as a source of technology transfer. They maintained that whatever they did had very little application, or I think Dr. Hecker coined the phrase first off, they were not very user friendly.

I frankly think it was a culture that was created by the very way they grew up, and also a culture on the part of the private sector that way—they did not need them. I mean we did not need anybody for a long time in the private sector. They did their own thing until the last 20 years.

But now I guess I am kind of amazed. Because essentially—and I am not talking about this in the perjorative or adverse sense—now we are all here talking about let us have the Federal Government take the lead in helping us get this done.

And I think that is great. It is very exciting.

What has happened?

Maybe one of the two of you, or maybe Dr. Hecker could tell us, just quickly, I do not want a long explanation, because Senator Bingaman has to leave, and I do, too.

Mr. ROLLWAGEN. Well, I will just start.

First of all, you are right. Perhaps, at least our company is not typical, in that we have enjoyed a close relationship with the national labs all along, and we have benefited greatly from that.

But I think, beyond that, we are in new territory or new water. We have not been here before. The kinds of technical challenges we are taking on in the industry as well as in government laboratories has reached a level of sophistication that strains our resources, and strains them both economically and intellectually, if I can put it that way.

And I think we all recognize that there are tremendous resources in the national laboratories, and that we can work together to solve some of these problems.

Secondly, I think that, particularly in the area of supercomputing, there is an infrastructure issue that literally is analogous to transportation systems—to the highway systems in the country. It is a chicken and egg problem. We need to get an infrastructure in place that makes this technology available.

It is very expensive. We are all going to benefit from it. And in the long run, I think it pays for itself and could be supported by user fees very directly. But it is a process of, how do we get this started? No one single company is going to step up and do it.

So I think those two factors may help explain at least why now is different from before.

Senator DOMENICI. Do you have a comment on that?

Mr. PALUCK. Yes.

I think the big thing that is new is the threat. Years ago the industry could stand alone and grow at its own rate and lead and have true leadership.

Today, our competition is fast becoming the greatest economic power in the world. And that is who we are fighting. And so we need all the help we can get. We cannot be small and stand alone. We need a catalyst to help. Not subsidies, but a catalyst to aid.

Senator DOMENICI. Dr. Connolly, do you have a thought on that?

Dr. CONNOLLY. Could you tell me the specific question again, Senator?

Senator DOMENICI. What is amazing to me is that when we started in technology transfer activities it was four or five years ago on a real serious basis, what is the hang up?

Why does it take so long to get agreements between the private sector and an industry? Why have the patents been held up so long? What about proprietary rights to things that are not patentable?

We had a chief scientist for General Electric, a vice president, writing a national thesis, published all over the country, his first name was Peter, I do not remember his last name, saying that national laboratories are not user friendly. We do not even need them. They do not respond to our needs.

And what I am saying is now it seems like we are having an awful lot of people, at least thus far, saying, as far as this activity, we want the government to take the lead and we want to work with them.

I just wanted a thought as to what is different, and I think I understand it. I think, frankly, it is what you both said, plus the area of private sector activity in this field to this point is minimal. And as a consequence, we do not have a whole batch of experience, pro or con, but rather, those who were involved are very concerned about using what they know about and getting the next benefit from it.

And we do not have experience in the marketplace yet that says governmental activities are rather frustrating. We will get that in about five years, after we have got it, but it will be solvable.

That was essentially it.

Dr. CONNOLLY. Yes, one of the difficulties in the computing industry is that it is hard to think long range because of your stockholders and things like that. You are very much concerned with the next quarter's profits.

Whereas the computing industry, in order to stay competitive, you really have to be thinking five to 10 years into the future. The Federal Government can do that. The computing industry finds it very difficult.

Senator DOMENICI. I just have two remaining questions, and they are both directed to Dr. Hecker.

Dr. Hecker, I am concerned and want to try to resolve this concern if it is real.

With the experience that we have had with something like the high-temperature superconducting center at Los Alamos, and we have two others in the country, in that one of the government's institutions, and we are going to have many of them involved in this network, had the capacity to put together this center, went out under a formal but not statutorily created charge from the Department of Energy in an effort to put it together, and is still plagued, but for the most current budget, the one that we are going to appropriate this year, with no funding for that effort because it is assumed that the institution, called Los Alamos Scientific National Laboratory, will fund it.

And you get an appropriation, and everybody is looking at your expenditure of monies up here, and occasionally somebody gets a wild idea, and I asked the GAO to see if they can review it and

make sure you are spending it right. And somebody finds superconducting—where in the world did they get that charge?

So I guess I am concerned that, if we are going to have institutions do collaborative work outside their normal one that is synergized in this, should we not be concerned about making sure there is something to be funded for them, whether it is NASA or you or whoever, in a very specific way to accomplish this mission? And is that why you have the collaborative R&D centers? Is that one of the reasons?

Dr. HECKER. It most certainly is.

Let me first set straight the situation on the superconductivity pilot centers. The way that the concept was developed, we had discussed with the Department of Energy funding those at $5 million per year, per center. And the Department of Energy did actually provide some seed funds, officially, for those centers.

So it was not that they were without Department of Energy funding. So it made it easier to get them going. Of course, we did use our abilities to have some internal discretion to be able to help those centers along. And in fiscal year 1991, they are in the President's budget for $5 million. So we had some monies.

In the area of the collaborative R&D centers that we are proposing for computing, I think there it is certainly true that we have relied principally on the fact that the nuclear weapons program has required the greatest computational capability in the world to be able to fund our end of the collaborations with industry.

And what we are saying now is that role should be formalized in such a fashion that we can do that in a variety of areas, especially at a time when the overall funding is decreasing substantially.

Senator DOMENICI. I have spoken to the Chairman, and he indicated that he would like to work with me and others on the committee on that concept. And we may be getting back to you for some information.

My last question, perhaps both you, Dr. Hecker, and you, Dr. Connolly, could answer for the record.

I am interested in what you think the structure should be to oversee this. Dr. Connolly, you have indicated that you do not like the way it is dictated in the bill. I am genuinely interested in making sure that it is managed but that it is flexible enough for the individual major institutions of the United States Government to do their own thing, but that we are in fact meeting some goals and objectives, and there is some direction.

And I would like your thoughts on that if you could give them to us in writing. And, Dr. Hecker, I would like yours in that regard. I do not think we have to over-manage this from the offset.

But let me close by saying, one of the most difficult things about this kind of approach is evidenced by the fact, Senator Bingaman, that we have one bill in another committee, a bill in this committee, and I understand there is yet another piece that has to go to armed services, maybe in a freestanding bill, what they can consider, and then we merge it up, and pass a bill.

That is extremely cumbersome, and yet, with good leadership, we might get it done. But I think we ought to be watching, when that piece is finished, how are we going to assure the maximum atten-

tion to funding it. Because now it is going to be a very mixed up, from the standpoint of our procedures around here, entity.

Is it funded in appropriations by the energy and water subcommittee, out of Department of Energy funds? Is it funded in appropriations by defense? Or is it funded by Senator Hollings over in commerce?

And frankly, I do not think this is an insignificant question. Because we could end up with a neat thing, then if it takes $60 million or $70 million or $80 million a year, aside and apart from the institutions that are spending their money, we could end up with it being nobody's baby. And that does bother me, because I have seen a couple of those happen up here.

I do really compliment Senator Gore for moving it. And I think we might move it through three or four committees up here and get it done. I wish it was not that way, but we might.

And I think we have to concern ourselves about that. And who you put in charge at the top might indicate the strength of the funding request, and also, to some extent, how successful we will be in getting it funded. That is, the general funding for the enterprise that we are calling a network, not for DOE's individual work or NASA's or someone else's.

So that is the reason I think we ought to be concerned about that top end.

My last observation: With reference to the private sector, and, Mr. Paluck, you are absolutely right, we have to drive this with goals that are achievable and measurable, not just research. And I think that, too, is very important that we decide how it gets managed, so that those kind of activities will be appropriately set, and somebody will pick and choose in a manner that is consistent with real needs and real capabilities, rather than the strengths of some agencies.

And that is a tough one, too. We will take a look at that.

Let me just close by saying that I personally appreciate all of your effort in this regard. And I think we are moving in the right direction. I am not going to be critical of Dr. Bromley. I think I know exactly what has happened. But I do think that if we get started, we are all going to be very fortunate that if Dr. Bromley and his new team are around for a while, they are exactly the kind of people that understand this.

The fact that they have worked this hard on it in putting all these agencies together, is a far cry from where we would have been two years ago or the year before, when we did not even have the Department of Energy considered or taken to the White House, when we are considering these new centers within the National Science Foundation that President Reagan announced.

We did not even have the Department of Energy in the package, nor were they ever invited to a meeting. And they only had about 40 percent of the scientists and engineers in the country, and were not even at the meeting.

So I think his role will be helpful once we get it going.

Thank you very much. It is good to be with you.

Senator BINGAMAN. Let me ask just a couple or three questions.

Let me allude, Mr. Rollwagen, to your chart. I think it is a very interesting chart, on the last page of your testimony, where you

compare the number of supercomputers in the United States and Japan by industry or by application.

It shows that the U.S. has 173 and the Japanese have 104, according to this chart. I think it is interesting that if you take out aerospace and petroleum and government, we have 48 and they have 92. So when you get away from the defense-related activity and petroleum, it is clear to me that they have a much greater use or application of supercomputers in the broader range of their industry than we do in the United States.

Is that a fair statement?

Mr. ROLLWAGEN. That is absolutely a fair statement. And it is noteworthy, at least, to see the difference in the market evolution in the United States versus in Japan.

In the United States, the market for these machines started exclusively with the government and has moved into the commercial sector, but still that represents less than half. Whereas the real market in Japan has been primarily in the private sector—exclusively in the private sector as far as our equipment is concerned. And it is a different orientation altogether.

Senator BINGAMAN. Dr. Hecker, picking up on that same point. These collaborative R&D centers that you are urging be established, it seems to me that one of the main reasons that that makes sense and really is essential is because of this lopsided situation we have, where the bulk of supercomputers in this country is in government. I do not think it is realistic to just say, okay, U.S. industry needs to do better.

If we are going to leverage what we already have going in our laboratories and in the federal sector and use that for the benefit of the private sector, we have got to find some new and innovative ways to do that. And I guess that is what I am seeing as one of the main justifications for your proposal.

Is that reasonable?

Dr. HECKER. It certainly is.

Senator BINGAMAN. Let me just also, Mr. Connolly, ask you this.

Senator Domenici indicated he would like something in writing, and, obviously, I think we all would, as to how to structure this thing. I compliment Dr. Bromley and I compliment Al Gore and all the rest of the folks who have been involved in trying to put this together.

But it reminds me very much of the famous statement that is attributed to Charles Lindbergh after he flew over the ocean. Someone complimented him for having been the first person to fly over by himself. He said, it would have been tougher if I had done it by committee.

And that is what we are trying to do here. We are trying to say, okay, we have had a good committee study of this problem. And they have laid out a course of action.

And now, since we do not have any executive branch leadership on the subject at all, why do not we try to have that same committee that did the study actually implement it? I am supporting the bill. But I am skeptical, as I gather you are, from your written comments and also your oral comments, that we can get from here to there with a committee in charge.

Now I mean obviously we can coordinate, but as far as actually taking the kind of decisive and significant action that is going to be required, we are going to have to have somebody in charge. And we are going to have to have some level of commitment for funding out of the executive branch of government.

Is that statement realistic?

Dr. CONNOLLY. I entirely agree with you.

The fix-it committee is very good at doing studies and making policy decisions and guiding the budget process. But to run a very large effort like this requires a kind of day-to-day flexibility, an ability to sort of say, well, hey, this approach does not work, let us try this one.

You are out there on the frontier of computing, you are going to run into a lot of glitches and bugs that are going to have to be dealt with somehow.

So you are going to need, and what I personally favor is some kind of third-party entity, perhaps a consortium of industry, university and government, with a strong management structure that would have some muscle to be able to coordinate the different agencies' efforts in this area.

And OSTP does a wonderful leadership job, but I do not think its committee structure can handle it.

Senator BINGAMAN. If any of the rest of you have suggestions on how this could be better structured--obviously, I think this was designed in order to deal with the political reality, which is, we do not have a commitment by the Administration to pursue this goal. And if we did, we would have a statute or a proposed bill in front of us saying the Secretary of Energy is in charge, or somebody is in charge of making this happen.

If any of you have thoughts as to ways to improve the management of this, I would be anxious to see those.

Again, I will thank everybody on behalf of what is left of the committee here. And I am sure that the committee will be acting favorably on the bill here in the near future.

Thank you very much.

[Whereupon, at 5 p.m., the hearing was adjourned.]

APPENDIXES

APPENDIX I

RESPONSES TO ADDITIONAL QUESTIONS

Dr. Sidney Karin
Director
(619) 534-5075

SAN DIEGO SUPERCOMPUTER CENTER

GENERAL ATOMICS, P.O. BOX 85608, SAN DIEGO, CA 92138-5608

April 9, 1990
SK-90-019

Hon. Wendell H. Ford, Chairman
Subcommittee on Energy Research
and Development
Senate Committee on Energy and
Natural Resources
Washington, D.C. 20510

Dear Chairman Ford:

Thank you for the opportunity to respond to your questions and those of Senator Domenici regarding S. 1976, the *Department of Energy High Performance Computing Act of 1989*. My responses are attached. I appreciate the difficulties of developing successful science policy, and I commend you and the members of your subcommittee and its staff for your concern and devotion to this important issue.

As a representative of a National Science Foundation supercomputer center that owes a great deal to developments by, and cooperation with, the DOE computer divisions at LLNL and LANL, I sincerely appreciated your invitation to testify before your Subcommittee on the DOE legislation. I regret that I was unable to be there in person because of an injury.

Please don't hesitate to call me if you require any additional information.

Sincerely,

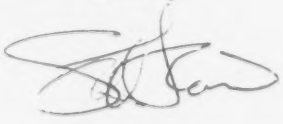

SK:Sjs
Enclosure

LOCATED at University of California, San Diego and administered by General Atomics.
CONSORTIUM MEMBERS: The Agouron Institute • California Institute of Technology • National Optical Astronomy Observatories • Research Institute of Scripps Clinic • Salk Institute for Biological Studies • San Diego State University • Scripps Institution of Oceanography (UCSD) • Southwest Fisheries Center • Stanford University • University of California at Berkeley • University of California at Davis • University of California at Irvine • University of California at Los Angeles • University of California at Riverside • University of California at San Diego • University of California at San Francisco • University of California at Santa Barbara • University of California at Santa Cruz • University of Hawaii • University of Maryland • University of Michigan • University of Southern California • University of Utah • University of Washington • University of Wisconsin

Sidney Karin, Ph.D.
Director, San Diego Supercomputer Center
and
Vice President, Advanced Computing
General Atomics

I. Responses to Questions from Senator Ford

Question 1: Please comment on whether there should be a lead agency or body on this overall effort, and if so, who that should be.

Answer: Having one agency provide the lead would be helpful to achieving the long-term objectives of legislation on the Federal High Performance Computing (HPC) initiative. The key agencies involved--DOE, DOD, DARPA, NASA, NSF, and NOAA--each has considerable experience and expertise in HPC, and each has made substantial contributions, i.e., research, software, service or other aspects of HPC. Any of these agencies could serve as an effective lead agency if appropriately chartered to do so. I ask only that we do not create yet another agency.

Although the NSF is responsible for managing the backbone network, NSFnet, no one is serving as the lead agency for HPC or, indeed, even for the Internet. The lead agency concept needs to be strengthened if it is to be effective prior to privatization. Whichever agency is chosen as lead, it needs to have both the responsibility and authority necessary to successfully execute its charter. Our key concerns here are that the production network be privatized as soon as possible and that barriers to interagency cooperation be removed. While there have been examples of cooperative efforts across agency boundaries, they have tended to happen despite bureaucratic impediments and to be relatively small scale. It would be desirable to provide up front incentives for cross agency activities, including large scale projects. Projects where one agency relies on another for critical path items and where the supplier agency has the production of such items in its charter are required to maximize the impact of federal funding, avoiding fragmentation and redundant efforts.

Question 2: Does S. 1976 place enough emphasis on education? If not, what should S. 1976 emphasize?

Answer: Yes, S. 1976 does strongly emphasize education, and we support the goals identified in (A) (B) and (D) of Sec. 7. (a) (5) in the bill. These are worthy goals, i.e.,

> (A) making the national laboratories available to senior graduate students, postdoctoral fellows, and faculty from the Nation's universities;
>
> (B) expanding summer science programs for high school students;
>
> (D) establish more cooperative research programs with the academic computational science community.

Since these goals are identical to the mission of the NSF Supercomputer Centers, we recommend a cooperative DOE/NSF effort, working through the existing NSF Supercomputer Centers.

The NSF Supercomputer Centers support computational researchers across a broad array of scientific disciplines. In addition, the NSF Supercomputer Centers are all located on the campuses of major universities; the infrastructure and the experience to assist academic researchers with education and training is in place and functioning well. Support for these efforts from DOE through the National Laboratories would be welcome.

Question 3: Should S. 1976 require more or less access to Department of Energy supercomputers by universities? Industry? Would the Collaborative R&D Centers Dr. Hecker suggests give sufficient access?

Answer: As I've noted, the infrastructure and experience to assist academic researchers and industrial users is in place at the NSF Supercomputer Centers.

The NSF Supercomputer Centers now supply broad access to advanced scientific computing facilities for academic researchers and modest access for industrial researchers; most of those accessing the NSF Centers are not associated with the mission-oriented DOE, DOD, and NASA facilities. We believe that the DOE should address the basic research community, within its areas of technology, in cooperation with the existing NSF Supercomputing Centers.

Some collaborative programs among federal labs, universities, and industry, which Dr. Hecker describes as Collaborative R&D Centers (CRDCs), are already coalescing. One example that Dr. Hecker mentions is the Gigabit Test Bed, an 800 Mbit/second network linking Los Alamos, Caltech, the Jet Propulsion Laboratory, and my own institution, the San Diego Supercomputer Center. What we need is for the Federal government to provide the funding to make these types of collaborative ventures successful; we don't need to create new Federal entities. We applaud Dr. Hecker's vision in recognizing that the Federal government must work together with universities and industry to optimize our limited resources and speed technology transfer. The DOE's additional support of and collaboration with the NSF Centers to further enhance their activities would significantly broaden the collaborative programs already in place within the academic and industrial communities. DOE Centers, particularly LANL and LLNL, have clearly contributed much already to the success of the San Diego Supercomputer Center. We support the goal of the DOE committing "substantial resources" to enhance greater Federal, academic, and industry collaboration. A collaborative DOE/NSF initiative would add Dr. Hecker's "substantial leverage" to the United States' long-term ability to advance new technologies.

Question 4: Some of the witnesses testified that S. 1976 could prevent flexibility by locking in specific expenditures within federal agencies. How should the bill be changed so as not to do this?

Answer: Philosophically, I would agree that computing technology is advancing so rapidly that it is difficult to determine what the appropriate balance of funding for networks, hardware, software, and research should be over the next five years. As the Committee knows, the Federal

government has undertaken numerous studies and, subsequently, made numerous recommendations on levels of funding. Recommendations differ, but all seek to assure that we provide the U.S. scientific and engineering communities with access to the most advanced computing resources available. And I must also add, to date, no part of the Federal government has committed the required resources.

I believe that the studies commissioned by the Congress and others in the early 1980s came to the correct conclusions and presented wise, practical recommendations. For example, the 1983 Bardon/Curtis report, "A National Computing Environment for Academic Research," was a particularly fine study and made balanced recommendations for upgrading the entire national infrastructure for scientific computing. I agree with that report. It recommended budgeting $110 million per year by 1986 for a few National Centers, and a similar amount for improving local facilities at universities.

And in light of these recommendations, the FY-1991 Budget Request for the NSF Centers, after nearly five years of existence, is roughly $60 million--significantly less than the Bardon/Curtis recommendation of $110 million (or the sum that $110 million would have become eight years later). I must also note that the funding level for the NSF Division for Advanced Scientific Computing for four Centers is roughly the same level of funding that NASA commits to its Numerical Aerodynamic Simulation Facility at the Ames Laboratory.

Panels and committees have studied the situation more than enough, but we have not achieved their recommended results. Therefore, it appears that it will take the foresight of the Congress to finally mandate sufficient funding for the Federal government's HPC programs. Strong, sustained, funding is necessary and a five-year authorization bill is important. And, of course, we all know that we must ensure that the subsequent appropriations bills are passed as well. The NSF Centers, with our negotiated five-year Cooperative Agreements, are perfect examples of the need to enact multi-year funding bills. The capriciousness of funding levels for the NSF has been destabilizing for the Centers, resulting in cuts to our negotiated Cooperative Agreements.

II. Responses to Questions from Senator Domenici

Question 1: In the evolution of TECHNET, the network I helped to start in New Mexico, the task of identifying database users was critical to its success. Could you elaborate on what steps your respective institutions have taken in identifying users, other than academic institutions, for example, chambers of commerce, companies large and small, agencies from the private sector?

Answer: As an NSF Center, our mission is to assure that HPC is incorporated into the broad U.S. research community. We are continuing to identify new non-academic users and introduce them to our intellectual environment. We are working with several dozen industrial organizations as well as the State government to help them work with all that our supercomputing environment has to offer, including our nationally recognized technical staff; advanced scientific visualization capabilities; a wide variety of applications programs; powerful computational systems; convenient login as well as access to national databases.

We offer the advantage of access to a leading-edge computational environment--opportunity to become immersed in our environment and work side by side with our scientists and engineers.

Question 2: Some of my colleagues who support this legislation in its present form, suggest that "hooking-up" institutions like yours will, itself, precipitate the educational process. This may be true for those like yourselves who already possess the necessary background and expertise, but could you tell me how others, who possess less background, will "interface" with this network? What would be needed for anyone owning a PC to be able to use this network?

Answer: The NSF Centers provide leadership, whose tangible manifestation is the training, education, and support required to assist those with less expertise and bring them "up to speed." We have already provided substantial assistance to the various state and local government sponsored supercomputer centers, as well as more than 100 universities

and dozens of industrial firms.

Indeed, a National Research and Education Network (NREN) will further enhance this nation's educational system. And our local regional network, CERFnet (California Education and Research Federation Network) is currently focusing on initiatives towards that end, as are most of the operators of other regional networks. CERFnet includes both educational and industrial members.

As far as how easy is it for the average person to access the network--anyone with a terminal or a PC can utilize the existing networks. Individuals can use dial-up service--you don't even need to have an expensive connection to be part of the Internet anymore.

The necessary mechanisms to train individuals in the use of the network and its resources are in place and functioning through such institutions as the regional network components of the Internet and the NSF Supercomputer Centers. What is critically needed now is expansion of the activities of the network engineering design and manufacturing and other commercial activities. The current arrangement continues to separate research activities from production activities, exacerbating our difficulties to compete in international markets.

Question 3: One approach to a supercomputing network has been attempted by the National Science Foundation. It has 5 centers that constitute "nodes." Would you comment on the benefits and problems associated with this approach and contrast it to one which placed the hardware at one facility.

Answer: NSF now supports four national supercomputing centers. Their facilities are indeed connected, as nodes on a wideband network called NSFnet, but it is very important not to confuse *centers* with *network nodes*.

A national supercomputing network is an important part of an overall strategy for high-performance computing. And supercomputing hardware of all kinds is essential, obviously. But the success of high-performance computing in advancing American competitiveness depends on more than

hardware and connectivity alone: it requires scientific leadership, education, and research--the various other components of S. 1976. These elements are united in the several NSF Centers, which are much more than "nodes" on a network.

The NSF Centers were established as scientific research institutions, intellectual centers with their own research and development programs and their own scientific staff to respond to the concerns of their peers and to assure that the computational environment for their academic disciplines remain state-of-the-art. These Centers have been immensely successful precisely because they don't just deliver "computer cycles" to the research community: they are part of that community; they conduct research and development programs and deliver the computational advances that community demands, and they provide *leadership* to the computational sectors of many scientific disciplines as well as the rest of the computing community.

The number of centers is about right. That is, the research community benefits from the friendly competition of a few centers, all seeking to provide the best in computation, communication, and visualization. Science and engineering require some redundancy, and a wise cognitive division of labor will differentiate among the resources, look at the best ideas, and also into other ideas at the same time. It's the principle of hedging bets and, as I noted in my testimony, it's a proper evolutionary strategy. Nature uses it, too. So, the right number of centers is a few--certainly not one in every state (the economic advantages of scale will be lost) or a single one or two facilities for everyone.

A single facility poses the same dangers as a single-party political system. It is placing all the eggs in one basket. No matter where it was sited, some natural disaster could occur that would put researchers all over the country out of action. It would be a prime target for sabotage, an unnatural disaster. It would eliminate redundancy. A single bureaucracy controlling access would be another kind of unnatural disaster, a guaranteed brouhaha bonanza. Please put the thought out of your mind.

Question 4: Could you describe for me the present situation regarding the accessibility of the supercomputing centers at

DOE facilities?

Answer: I do not know how the classified facilities arrange access. I do know that DOE gives access to NMFECC, which is an open facility in which no classified work is done, to researchers funded by DOE, whether they are located at DOE laboratories or at universities. The university researchers can reach NMFECC via the Internet. MFEnet and HEPnet are now migrating to a single dedicated network, called ESNET, with nodes at fusion research facilities and a number of universities and laboratories. Access is provided to those whose work is related to the DOE mission. ESNET is linked to the NSFnet as part of the overall Internet. LANL is also linked to SDSC and Internet, too.

Question 5: In your opinion, what does this legislation need to properly address the question of network management. Who (or what) in your opinion is best suited for this?

Answer: The production network should be privatized as soon as possible. The NSFnet has operated extremely well under the management of MERIT, along with IBM and MCI, and may be a useful model for managing the NREN.

As stated in my response to Senator Ford's first question, the lead agency concept needs to be strengthened if it is to be effective prior to privatization. Whichever agency is chosen as lead, it needs to have both the responsibility and authority necessary to successfully execute its charter. Much of the academic community will need subsidized access over the national network if we are to succeed in fostering its widespread use.

28-910 O - 90 - 6

4/7/90

RESPONSES TO
SENATOR FORD'S QUESTIONS FOR THE RECORD
ON S. 1976
FROM
S. S. HECKER
LOS ALAMOS NATIONAL LABORATORY

Question #1:

Please comment on whether there should be a lead agency or body on this overall effort and, if so, who that should be.

Answer #1:

High-performance computing (HPC) is of such pervasive importance to the United States that strong leadership in this area is essential. We see such leadership coming from several, rather than a single, government agencies. These include agencies that have critical HPC requirements (such as DOE, DoD, DARPA, NSA, NASA, and NSF) and that have extensive experience in high-performance computing.

Although many of the HPC requirements are generic, there are a sufficient number of agency-specific requirements that no single agency should play the lead role.

There is, however, an urgent need for coordination for several reasons. First, the generic needs of these agencies should be coordinated. Second, there are a number of other federal agencies with emerging HPC requirements (EPA, NIH, DOI, Education, etc.), and central coordination will help them benefit from the work of the current major HPC agencies. Third, the needs of U.S. industry are not adequately covered by any single agency, and hence central coordination will benefit the industrial requirements.

We believe that the Office of Science and Technology Policy (OSTP) is in the best position to provide overall coordination and guidance.

Question #2:

Does S. 1976 place enough emphasis on education? If not, what should S. 1976 emphasize?

Answer #2:

As in all other areas of science and engineering, more emphasis can always be placed on education - starting with the elementary school level and continuing to higher education. However, one of the four major components of the OSTP High Performance Computing Program along with roughly 25% of the resources of the Program are directed at education and basic research. Similarly, S. 1976 follows the lead of the HPCP in both intent and resource allocation (assuming some of the resources of the other three major components is actually used for educational activities). So education is in fact being emphasized, and the level of emphasis is appropriately balanced with the rest of the Act.

During the 1980's, the United States acquired almost 50 million personal computers. Many of the users are students. Over time, the age and grade level of these students have continued to become lower. The time is right to introduce them to the next level of high performance computing.

The HPC program and S. 1976 could emphasize more specific programs to support access of students to HPC environments. Our Collaborative R&D Center (CRDC) proposal has a strong educational component, both at the high school and university level. At the university level we would sponsor graduate student research on HPC at universities and offer graduate students the opportunity to work at the CRDCs to gain hands-on experience.

Question #3:

Should S. 1976 require more or less access to Department of Energy supercomputers by universities? Industry? Would the Collaborative R&D Centers you suggest give sufficient access?

Answer #3:

Department of Energy policy has been to provide access to the DOE supercomputers by universities and industry. However, at your March 6, 1990, hearing we heard university officials call for increased access.

The Collaborative R&D Centers would go beyond allowing access; they would encourage it by building it into the CRDC concept. The CRDC concept calls for distributed centers with multiple participating organizations. Universities would be included as CRDC participants. In addition, the university collaborations objective of the CRDC would specifically fund university research and graduate student participation.

Industry access is also crucial. Again our belief is that we must provide more than access. We must work collaboratively with industry to help demonstrate the benefits of HPC in their business operations, for inventory control, for design, and for manufacturing. Again, the CRDCs are designed to do just that through an industrial user program.

Question #4:

Some of the witnesses testified that S. 1976 could prevent flexibility by locking in specific expenditures within federal agencies. How should the bill be changed so as not to do this?

Answer #4:

Flexibility is a key requirement for all R&D funding. It is especially important in rapidly moving fields such as HPC. Hence, any attempt at this point to lock in funding for specific technical approaches in HPC for the next five years would be counterproductive.

However, there is no doubt the HPC is a strategic technology for U.S. military and economic strength, as well as significantly benefiting the quality of life. Hence, enhanced funding such as that outlined in the Bromley OSTP High Performance Computing Program is essential. Each of the major agencies should develop a plan for an

accelerated program in HPC. Funding for such programs, with specific technical approaches being re-evaluated from time to time and not being locked into specific technical approaches over a long period, should begin in FY1991.

Question #5:

Please elaborate on the concept of Collaborative R&D Centers you mentioned in your testimony. Where would the centers be located? How much would the centers cost? Should the legislation specify a certain number of these centers? What would the centers do? How would the centers operate? How do they compare to the National Science Foundation's supercomputer centers? Should the centers always have the most powerful computers?

Answer #5:

I propose developing Collaborative R&D Centers (CRDCs) designed to keep the United States at the leading edge of computing performance through innovative collaborations among industry (both computing technology vendors and industrial users of computing technology), the federal laboratories, and universities. The concept of the CRDC is developed further in the accompanying Los Alamos white paper on CRDCs.

The CRDCs would have goals in direct support of the intents of the HPCP and S. 1976 by (1) helping shape the leading edge of high-performance computing and (2) by putting more useful computing power into the hands of U.S. industry and broader segments of the population.

The most important feature of a CRDC would be (1) collaborations with the computing vendors, designed to keep the United States at the leading edge. Specific activities will include:

- Challenge computational performance and tools with demanding applications.
- Identify requirements for computing technology developers and manufacturers.
- Help define, develop, and test the future computational environment, including integrating new architectures, software, and visualization techniques.
- Develop software, hardware, algorithms, etc., in collaboration with computing vendors.
- Help develop and test new distributed networking systems.
- Disseminate knowledge and experience gained to other users.

Other important features of a CRDC will include (2) university collaborations designed to provide new, innovative ideas in computational science and educational opportunities for graduate students; (3) an industrial user program to encourage increased use of analysis and modeling with supercomputers in U.S. business and industry; (4) a general user program to provide user-friendly computing through networks to local communities for business, schools, and universities; and (5) developing enabling technologies for high-performance computing such as free-electron laser lithography, superconducting electronics, and opto-electronics through laboratory-industry collaborations.

In my opinion, the most effective structure of a CRDC is not an entity associated with a single organization but rather a geographically-distributed consortium of participants (such as universities and laboratories) led by an organization that has very strong credentials in high-performance computing.

It is envisioned for the out years that there would be 2-3 CRDCs led by DOE laboratories

and a total of 3-5 CRDCs nationally; there would be a smaller number in the beginning of the HPCP. Federal funding of a CRDC would be required on the order of $30M per CRDC per year initially and growing to the order of $40-50M per year per CRDC in the out years. Industrial participants in a CRDC would pay for their own involvement, although a CRDC could fund the research and development associated with some prototypical computing technology.

The mission of the NSF supercomputer centers is to provide, via networks from the centers, regional supercomputer resources to serve regional educational institutions, including a strong research component in the center itself. The CRDCs will have to have all of these attributes plus several others: the interaction with industry vendors, the interaction with industry users, and a strong in-house capability in R&D aspects of the computing infrastructure to be able to really help industry users and vendors with their leading edge and other non-routine problems.

It is neither necessary nor appropriate for each CRDC to have a complete complement of supercomputers and experimental high-performance architectures. Hardware and software systems that complement the particular area of emphasis and expertise of each CRDC should be contained in that CRDC. The high-speed network infrastructure will make these capabilities available to interested parties at other CRDCs. CRDCs will serve as sites for vendors to test new ideas against real applications before their products are brought to market.

4/7/90

RESPONSES TO
SENATOR DOMENICI'S QUESTIONS FOR THE RECORD
ON S. 1976
FROM
S. S. HECKER
LOS ALAMOS NATIONAL LABORATORY

Question #1:

In the evolution of TECHNET, the network I helped to start in New Mexico, the task of identifying database users was critical to its success. Could you elaborate on what steps your respective institutions have taken in identifying users, other than academic institutions, for example, chambers of commerce, companies large and small, agencies from the private sector?

Answer #1:

The integrated computing network at the Los Alamos National Laboratory currently has more than 8000 clients (or users) nationwide. Of these users, about 5000 are Los Alamos employees and about 3000 are employees of other organizations. The user community at Los Alamos has evolved with the programmatic mission of the Laboratory as the major driver. The National Laboratories are prevented by law from competing with private industry to provide data center or networking services. Thus each external organization (there are currently about 50) that uses Los Alamos computing or networking services has required the explicit approval of the Department of Energy. Such approval is given when it can be shown that the support to the external organization can be provided without interference to DOE programs and that the Laboratory provides a unique service that is not available commercially. Examples of such services are remote classified high-performance computing and unique software available only at the Laboratory. Additionally, the Laboratory can provide computing and networking services to an external organization when mutual interest exists such as when the two organizations are working on the same project.

In this environment the Laboratory has provided its computing and networking capabilities within the constraints mentioned. A fairly large number of the external users of the Los Alamos facilities have resulted from the potential user making the initial contact inquiring about the availability of resources.

Question #2:

Some of my colleagues who support this legislation in its present form suggest "hooking up" institutions like yours will, itself, precipitate the educational process. This may be true for those like yourselves who already possess the necessary background and expertise, but could you tell me how others who possess less background will "interface" with this network? What would be needed for anyone owning a PC to be able to use this network?

Answer #2:

Network software and protocols are available to allow nongovernment sites to interface

with national networks. With sufficient training and documentation, it would be relatively easy to integrate these users with the network. A trained system and network administrator who can interact with others on the network will be necessary for each group of users. Documentation and preliminary classes to introduce users to the network and facilities should be fairly straightforward.

The interface would be as follows: A consulting office would be created and operated by the Collaborative R&D Centers (CRDC) that we have proposed. It would be accessible by phone, fax, drop-in visit, or electronic mail. It would provide information on setting up one's personal computer (PC) to interact with the network. It would provide software and installation instructions for the user to load onto his own PC. This software would have suitable help programs to guide the user in learning how to connect his PC to the cable or phone network, how to setup the appropriate modem, how to log on the system, what programs he will find on the system, and how to use these programs. It will also provide the user with information on tieing his PC to other commercial networks. The CRDC would also set up a suitable user-friendly program on the host computer, and provide guidance in its use. These services, offered to individuals, would also be available as training classes to organizations requesting presentations to their membership, students, or employees. The consulting office might also have a user PC setup in its office for one-on-one training of individuals.

The PC owner will need to acquire the following items:

- A PC (his expense, unless school or nonprofit outfit)
- A modem (ditto)
- Software package for his particular PC (modest or no cost)
- Training (via the above, or video tape checkout from public library)
- Passwords; log-on procedures; friendly interface to word processor, spreadsheet programs, data base software, high quality output, compilers, data networks, etc.; all from CRDC (modest or no cost, but probably at some user-borne cost above a certain minimum)

Question #3:

One approach to a supercomputing network has been attempted by the National Science Foundation. It has five centers that constitute "nodes." Would you comment on the benefits and problems associated with this approach and contrast it to one which places the hardware at one facility?

Answer #3:

The concept of NSF Supercomputing Centers has generally worked very well. The bottom line is that it has increased the availability of high-performance computing (HPC) to the university community and has also had some benefits in the private industry sector.

The question of the utility of nodes is much more one of considering the virtues of a few (such as the five at NSF) versus hundreds and not the five versus one as posed in the question. We believe strongly that there should only be a few centers that provide the top end of the computing spectrum because of the difficulty and expense associated with such a facility. This pertains not only to the hardware and vast associated infrastructure but also to the whole user environment. The benefits of a few powerful computing nodes have been demonstrated by the NSF Supercomputing Centers, the Los Alamos Central Computing Facility, and the DOE's Lawrence Livermore National Energy Research

Computation Center and its network.

We would like to add that in addition to having a few such "user" centers, we believe that the concept of Collaborative R&D Centers that we have proposed fills another important requirement, that is to help the United States stay at the leading edge of high performance computing and make each new advancement available to the university and industrial communities. Again, the CRDC concept calls for a "few" such centers.

Question #4:

Could you describe for me the present situation regarding the accessibility of the supercomputing centers at DOE facilities?

Answer #4:

Access to DOE supercomputer centers currently is available through a wide variety of federal and commercial networks using a variety of mechanisms such as dial-up telephone access, satellite links, and dedicated point-to-point communication channels such a fiber optic links.

Access is controlled in the sense that the DOE centers do not provide computing services in competition with commercial computing centers and services. Both the academic and industrial communities therefore do not have free access to DOE computers. In general, all external users of a DOE center have to be in collaboration with the institution operating the center on technical projects of mutual interest to both parties. Industrial users reimburse the DOE center for computing services provided.

In addition to the collaborative requirement of the work, access to DOE classified resources is limited to those agencies that have the required security approvals and to networks protected by National Security Agency approved encryption devices.

The DOE National Energy Research Computation Center at Livermore is a national center that supports users of the DOE energy research community in a sense similar to the NSF supercomputer centers. Potential users write technical proposals requesting computing time for their particular projects; if approved by a committee, an appropriate block of computing time (and associated resources) is granted to the user.

At Los Alamos in particular, both the academic and industrial communities use our Central Computing Facility in the sense outlined above. Blocks of computing time on our production machines (the Cray machines, for example) are not made available for any community of user to have unrestricted access to for any purpose other than for collaborative interactions with Los Alamos programs and research. The industrial users of our production machines reimburse us for the use of the computing services provided by us.

However, there are two advanced (non-production) computing centers [the Advanced Computing Laboratory (ACL) and Numerical Laboratory] at Los Alamos that promote the use and investigation of computers that are either of a prototype nature or novel architectures of a performance level approaching that of current supercomputers. The ACL is involved in several consortia (the DOE Office of Energy Research Advanced Computing Program, the NSF S&T Center for Research on Parallel Computation, and the proposed national gigabit testbed), each of which has a goal to provide a variety of state-of-the-art technology for use by researchers throughout the country. The research

on these computers spans a wide variety of areas such as scientific algorithms, visualization, and networking projects. These facilities are made available free of charge to qualified researchers to pursue independent research in a variety of fields. By submitting a proposal outlining the goals, benefits, and approach of the specific research project, qualified researchers may be allotted time on available machines. These facilities are available through national networks and are in an open, unclassified environment.

As part of the activities of a CRDC supporting the High Performance Computing Program, it would be the intention of Los Alamos to not only make broader access to our computers (both research and production computers) possible, we would collaborate with a broader user base in the context of helping them learn how to use supercomputers more efficiently for their own programs and not just for programs of mutual interest to Los Alamos.

Question #5:

In your opinion, what does this legislation need to properly address the question of network management? Who (or what) in your opinion is best suited to this?

Answer #5:

We recommend that the process being established by the Federal Network Council be used for the oversight of the development and management of the network developed for the HPC program. Initially, we do not believe that it is appropriate for a commercial organization to manage the network before the required research and development for the network is completed.

The HPC network protocols will need to be developed and defined by the agency responsible for networking all HPC resources. This should follow the TECHNET model exactly in that the protocols used by TECHNET were established standard protocols when selected by the TECHNET team. Once the required research and development for the NREN is completed it may then become appropriate to turn over the administration of the network to a commercial organization.

COLLABORATIVE R&D CENTERS FOR HIGH-PERFORMANCE COMPUTING

S. S. Hecker
Los Alamos National Laboratory

EXECUTIVE SUMMARY

As part of the implementation of the national High Performance Computing Program (HPCP) and as my recommendation in particular response to Senator Johnston's bill S. 1976 for DOE, I propose developing Collaborative R&D Centers (CRDCs) designed to keep the United States at the leading edge of computing performance through innovative collaborations among industry, the federal laboratories, and universities.

The CRDCs would directly support the goals of the HPCP and S. 1976 by (1) helping shape the leading edge of high-performance computing and (2) by putting more useful computing power into the hands of U.S. industry and broader segments of the population.

The most important feature of a CRDC would be (1) collaborations with computing vendors designed to keep the United States at the leading edge. Included are activities specific to:

- Challenge computational performance and tools with demanding applications.
- Identify requirements for computing technology developers and manufacturers.
- Help define, develop, and test the future computational environment, including integrating new architectures, software, and visualization techniques.
- Develop software, hardware, algorithms, etc., in collaboration with computing vendors.
- Help develop and test new distributed networking systems.
- Disseminate knowledge and experience gained to other users.

Other important features of a CRDC will include (2) university collaborations designed to provide new, innovative ideas in computational science and educational opportunities for graduate students; (3) an industrial user program to encourage increased use of analysis and modeling with supercomputers in U.S. business and industry; (4) the development of enabling technologies for high-performance computing such as free-electron laser lithography, superconducting electronics, and opto-electronics through laboratory-industry collaborations; and (5) a general user program to provide user-friendly computing through networks to local communities for business, schools, and universities. Each of the five areas will draw upon the academic community to bring added value and strength to the CRDC.

In my opinion, the most effective structure of a CRDC is not an entity associated with a single organization but rather a geographically-distributed consortium of participants (such as universities and laboratories) that is led by a lead organization that has very strong credentials in high-performance computing.

It is envisioned for the out years that there would be 2-3 CRDCs led by DOE laboratories and a total of 3-5 CRDCs nationally (including DOE's); there would be a smaller number in the beginning of the HPCP. Federal funding of a CRDC would be required on the order of $30M per Center per year initially and growing to the order of $40-50M per year per CRDC in the out years. Industrial participants in a Center would pay for their own involvement although a CRDC may fund the development of some prototypical technology.

INTRODUCTION

High-performance computing in the United States should be driven by two primary objectives: (1) we must stay at the leading edge of computing to solve increasingly demanding problems of science and technology, and (2) we must provide supercomputing power to greater segments of society to enhance new capabilities, products, and conveniences to business and the home. The first objective will keep a focus on useful computing performance and computing tools versus computing for computing's sake. The second objective is an affirmation that in its broadest definition national security (defense, economic competitiveness and productivity, and quality of life in general) are increasingly dependent on computing and can be strengthened by more effective use of computing.

The bottom line is that a country that seeks to control its future must lead in the development and application of high-performance computing.

The proposed High Performance Computing Program (HPCP) by the Office of Science and Technology Policy (OSTP) plus legislation in the form of Senator Johnston's bill S. 1976 and Senator Gore's bill S. 1067 to carry out the HPCP all have the intent to:

- Maintain and extend U.S. leadership in high performance computing and encourage U.S. sources of production
- Encourage innovation in high performance computing technologies by increasing their diffusion and assimilation into the U.S. science and engineering communities
- Support U.S. economic competitiveness and productivity through greater utilization of networked high performance computing in analysis, design, and manufacturing

As summarized in the OSTP Program that laid out a plan aimed at enhancing the competitiveness of the United States in high-performance computing technology, we need to develop a foundation using basic research and human-resource development. On that critical foundation the U.S. needs programs to (1) develop high-performance computing systems and advanced system software, (2) increase the resource-sharing capabilities across the nation by developing a National Research and Education Network (NREN), and (3) support Grand-Challenge research (as defined and summarized in the OSTP report) through the development of software components and tools and algorithm developments.

Many of the scientific Grand Challenges and many of the most demanding technical problems are either directly a government responsibility or are of intense interest to society. The key technology problems include those in defense, space, environment, health, and human welfare. Therefore, I support a strong role for the U.S. government to keep the nation at the leading edge of computing.

Yet for high performance computing to have its greatest impact there must be full collaboration with U.S. industry - the computing vendors as well as the computer users (including potential users). Together, government and the computing technology industry must keep the United States at the leading edge of computing technology and make user-friendly supercomputer power available to a greater fraction of industry, business, and the public.

Several government agencies have and continue to play an important role. The mission of the DOE in particular has always required the application of leading-edge high-performance computing: first, to problems of national security and now to challenges

of energy, the environment, and health. In pursuing its mission, the DOE has established in its national laboratories the world's most powerful scientific computational capability.

As a consequence, the DOE and its laboratories are in a position to help the U.S. maintain its leadership, foster the strengthening of the U.S. computing industry, and help disseminate high-performance computing into analysis, design, concurrent engineering, and manufacturing for U.S. industry. The laboratories are perhaps the most sophisticated, demanding users that push the requirements of high-performance computing to its limit; they have an unmatched record of close working relationships with computing vendors; and they have excelled in the development of high-performance computing technologies and computational science.

At Los Alamos, we have dedicated ourselves to the goals of a national computing initiative for over four decades. We have the environment and motivation to help lead an effort aimed at enhancing the development of high-performance computing technologies in cooperation with industry and universities.

My proposal for CRDCs to help support the goals of a national computing initiative is specifically in response to S. 1976, but it also supports the goals of the OSTP program in general. To encourage even greater interaction between the supercomputing users and suppliers, I propose developing a few (2 to 3) CRDCs led by DOE laboratories designed to keep the United States at the leading edge of computing performance and find solutions to important applications through innovative collaborations among industry, the federal laboratories, and universities.

LOS ALAMOS AND SUPERCOMPUTING: THE BASIS FOR A ROLE MODEL

Supercomputing originated with the "Los Alamos problem," as predicting the performance of the first atomic bomb was called during the Manhattan Project. The first attempts at supercomputing (a term relative to time) were made with arrays of mechanical desktop calculators organized by teams of operators. In 1945, Los Alamos researchers were the first demanding users of the University of Pennsylvania ENIAC, the world's first large-scale electronic computer.

Shortly after that, Los Alamos implemented the John von Neumann "stored program" architecture into its MANIAC computer, with which it ran the first large-scale hydrodynamic simulation of a thermonuclear weapon. The MANIAC was designed and built by Los Alamos scientists.

In 1953 IBM produced its first von Neumann machine with capabilities similar to those of the MANIAC, the IBM model 701. Serial number one was installed at Los Alamos. This began the era of Los Alamos collaboration with computing vendors that continues today.

The first IBM connection lasted into the mid 1960s, at which time Los Alamos began to work closely with Control Data Corporation to improve its machines, which, in short order, became the fastest in the world. In time, Los Alamos had eight of Control Data's supercomputers.

The first Cray-1 was taken to Los Alamos by Seymour Cray for testing and evaluation in 1976 when it was the only asset of the start-up company, Cray Research, Inc. The close collaboration between Cray Research (the entrepreneurial supplier) and Los Alamos (the sophisticated and demanding user) was key to the success of Cray supercomputers and

to the viability of the young company.

Today, Los Alamos is the most powerful scientific computing center in the world with the equivalent power of about sixty-five of the original Cray-1 computers serving more than 8,000 users throughout the nation via a national computing network. But it is our own computationally intense applications, starting with the complexity of designing nuclear weapons and now including a wide range of programs critical to emerging national and global challenges such as energy, ecological, and economic concerns, that require us to remain at the leading edge of computing technology.

The technology supporting supercomputing has grown rapidly. Over the past forty years, some applications have enjoyed an increase in computing speed of ten billion, or ten orders of magnitude. This increase has come in almost equal parts from advances in computer hardware and from more sophisticated numerical algorithms. In the past few years, it has become increasingly apparent that new algorithms and massively parallel computer architectures offer the greatest potential for dramatic performance improvements. In addition, with increased speed has come the recognition of other significant requirements for large-scale computing:

- Real-time visualization,
- High-speed networking to provide convenient and powerful computing capability at your desk,
- Software to facilitate writing sophisticated scientific computer programs

So, Los Alamos has turned much of its attention and collaborations in these directions. The high-performance computing environment we are developing consists of powerful workstations at an individual's workplace connected to high-bandwidth and functionally transparent networks that provide connectivity to supercomputers. Using this scheme, we seek to maximize the productivity of the individual using the integrated system. Hence, the task is one of integrating these different functionalities (workstations on desktops, networks, supercomputers) into a seamless environment that can be geographically distributed.

We have made progress toward the integration of these technologies. As a major step, we have established an Advanced Computing Laboratory (ACL) at Los Alamos. It was designed to be a catalyst for leading-edge research in these three key technologies. One goal of the ACL is to help develop an integrated computing environment from these separate pieces. A current focus of this laboratory is the second-generation Connection Machine. It has 65,536 processors, 2048 megabytes of memory, 4 frame buffers and monitors, and a 20-gigabyte data storage system.

The Advanced Computing Laboratory also serves to foster interactions with the national computational science community and is the primary interface for Laboratory participation in several team activities or consortia: the DOE/Applied Mathematical Sciences Advanced Computing program, the NSF Science and Technology Center for Research on Parallel Computation (a consortium of Rice University, Caltech, Los Alamos, and Argonne), and the Gigabit National Test Bed. Los Alamos' ACL is an excellent example of an existing center of excellence in high-performance computing that is positioned to participate in the development of computing technology through the High-Performance Computing Program. The activities of the Advanced Computing Laboratory are driven by the need to solve a broad spectrum of very computationally intense applications. I believe that the ACL is a prototype for the Collaborative Research and Development Centers for high-performance computing that I am proposing.

COLLABORATIVE R&D CENTERS

The HPCP needs to expand and build upon the Los Alamos model of extensive interactions and collaborations and experiences just described; it would be both too time consuming in this rapidly-changing technological environment and prohibitively expensive to start from scratch. The challenges of applications are increasing too rapidly, plus the applications and the computing tools needed for their solution are too big and too multidisciplinary for any one organization to be able to address effectively by itself. This model is what I refer to as a Collaborative R&D Center.

The Japanese can now build single processors as fast as any processor we have in the United States. The next few years definitely favor even more rapid Japanese advances because of their dominance in the memory chip market and in new materials such as gallium arsenide. Conversely, the United States will depend on better integration of processors into systems and systems into networks to strengthen our current lead in high-performance computing. It is crucial for the United States to enhance its integration advantage by encouraging and facilitating closer collaboration between demanding users and innovative suppliers.

A computing environment that is friendly and available to large segments of our population would likely change the entire fabric of American life. Business decisions and operations already depend largely on computing, but complete connectivity and access to national data bases with powerful computers would increase their productivity and hence the nation's international competitiveness. Access to data bases could influence daily decisions that vary from health care and diagnosis to automobile repair.

Supercomputers and the associated technologies have been developed for scientists who are traditionally willing (in fact, eager) to use the most powerful and appropriate tool that will allow them to solve their problems. This approach has resulted in a collection of capabilities that are powerful for solving problems but are not always considered to be easy to use by a less determined audience. If high-performance computing capabilities and concepts are to be integrated into the real world of business and leisure, a considerable amount of research and development will be needed to make these systems user friendly, so that the power of a supercomputer is accessible through the touch of a mouse on a desktop, either in the office or at home.

The greatest contribution that Los Alamos (and several other DOE labs) can make to the High-Performance Computing Program is to shape the leading edge of computing. In the past, we accomplished this role almost entirely with the financial support of the nuclear weapons research, development and testing program. That is no longer appropriate, nor possible, today. As mentioned above, our contributions now cover a broad spectrum from the human genome project to enhanced oil recovery. From this application base we wish to enhance the most dominant trends that we see in computing today: (1) the continued push toward faster computers, (2) the evolution toward the super workstation system, and (3) the development of national network technology. All will require the development of advanced software and algorithms to turn these trends into useful advances in computing performance.

We can support the HPCP and U.S. competitiveness by

- Helping to shape the leading edge of HPC, and
- Putting more useful computing power into the hands of U.S. industry and broader segments of the population.

These two goals can be effectively achieved by broadening the objectives of the "Los Alamos" role and by establishing several centers that are both *collaborative* and *distributed*.

The Distributed Concept

A CRDC in high-performance computing will be most effective as a distributed center comprised of a lead institution, several principal participants, and a number of affiliated organizations; a CRDC is in effect a consortium of several organizations. The lead institution must be an established high-performance computing center with a broad base in the technologies to be developed and must take responsibility for direction and performance of the CRDC. While individual sites will make important contributions, the geographically distributed nature of the national computing resource must be key in the organization of the CRDC itself. Thus, CRDC high-end use of networking and other technology to effect distributed computing will play an important role in the broader goal to develop efficient and distributed massively-parallel computing over a large geographical area.

The Collaborative Concept

The CRDCs will be collaborative by bringing together the resources and talents of (1) an institution like Los Alamos having a spectrum of Grand-Challenge applications supported by a major supercomputing center plus advanced research computing tools along with (2) universities to draw upon fresh new, innovative ideas and to provide an opportunity for training future computational scientists, (3) industrial users to engage their applications plus help industry take greater advantage of supercomputing, (4) other federal laboratories to take advantage of their ideas, users, and applications, and the (5) computing technology vendors to participate in synergistic development of computing tools in response to the drive for better solutions to applications.

These collaborations would be cost shared by the federal government and industry.

Major Activities of CRDCs

A CRDC led by a DOE laboratory would be part of the implementation of the High Performance Computing Program. Its strength would arise from the push of challenging applications helping to shape the leading edge of computing. The CRDC should complement other activities supporting the HPCP undertaken by other agencies.

Some of the components and activities of a CRDC would include:

1. *Collaborations with computing vendors* -

- *Challenge computer performance and computing tools with demanding applications.* In the same way that the DOE laboratories have provided the "pull" for high-performance computing technology from the computing industry in the past, the CRDCs would challenge high-performance computing vendors to continue pushing the state of the art.

- *Identify requirements for suppliers and inventors.* Just as the DOE laboratories have shown high-performance computing vendors the computing environment that will exist in the future, CRDCs, with their demanding users and

applications, would help high-performance computing vendors understand what products will be needed next. A significant role of the CRDCs would be to encourage growth of the distributed center network to include major industrial participation.

- *Test this new computational environment, integrating new architectures, software, and visualization.* The CRDCs must collaborate among themselves as a distributed center for testing the national computing environment. Several current examples of cooperative efforts are the DOE Applied Mathematical Sciences effort, the NSF Science and Technology Center for Research on Parallel Computation, and the National Gigabit Test Bed (in all of which Los Alamos is a major participant).

- *Develop software, hardware, algorithms, etc., in collaboration with high-performance computing vendors.* Software and hardware development often benefit immensely from collaborations. Some software, for example, has proven difficult or impossible for vendors to develop because they typically do not have the demanding applications in house. Innovative hardware ideas often originate with users that are closest to the problems. Vendors are then invited to collaborate in the developments, and the resulting products may become commercially viable. An example is the High Performance Parallel Interface cross-bar switch originally conceived at Los Alamos, prototyped with Digital Equipment Corp., and now marketed by Network Systems Corp.

- *Help develop and test new distributed networking systems.* The CRDCs will undoubtedly build or drive the development of high-performance computing environments consisting of hardware and software from many different vendors. These equipment and networking concepts will be tested on demanding applications.

- *Disseminate knowledge and experience gained to others.* The CRDCs will play a crucial role in education and technology transfer to other centers, to industrial vendors, and to industrial users, primarily through the vehicle of collaboration. Specific roles for CRDCs include: encouraging U.S. computing industry to stay at the leading edge of competitiveness and technology; providing direct help to U.S. industry in applying computing technology to their industrial needs; and helping industrial concerns who are currently not using advanced computing technology optimally in their research, development, and production processes.

2. *Academic collaborations* -

Basic research is a key to the success of high-performance computing. Vigorous university programs in computational science, electrical engineering, mathematics, materials science and computer science will provide long-term guidance and stability to the HPC effort. This research emphasis in the university community will be coupled to their educational responsibilities for raising a generation of computational scientists. It is important to provide training and support to teachers and educators, from elementary school level through postgraduate university level, so that their teaching can better stimulate and encourage the appropriate growth of human resources in the computtational sciences; and, offer direct training environments to students, again in support of schools and universities.

(a) It is imperative to involve university faculty to address more effectively the Grand Challenges and collaborate with the computer technology vendors in the activities

outlined in the first component above. The product of such collaborations would be new, innovative approaches to hardware, software, and simulation algorithms.

(b) Secondly, a very valuable component of a CRDC would be to specifically involve significant numbers of graduate students (say on the order of a 100 for a given CRDC) to provide a larger, more qualified future pool for the United States in computational science. The CRDC would provide fellowships to fund the research of the graduate students doing research in support of the above item 2(a). Furthermore, the Center would provide opportunities for the graduate students to work and study at the Center, interacting with all the participating organizations comprising the Center.

3. *Industrial user program* -

In John Rollwagen's statement at the March 6, 1990, hearing on S. 1976, he presented data indicating that industrial users of supercomputers in Japan outnumber comparable U.S. industrial supercomputer users by roughly a factor of two to one (except for aerospace where the United States clearly has the lead). Such a comparison would indicate that most U.S. industry is not using supercomputing as effectively as it could be.

Therefore, a CRDC would:

(a) Encourage U.S. industry use of analysis and modeling within an HPC environment (provide CRDC support to industry not only in computing but also in analysis and modeling).

(b) Develop and make available to U.S. industry user-friendly HPC capabilities through high-speed networks. Design, manufacturing and business applications would be targeted.

4. *Enabling technologies for high-performance computing* -

Other agencies such as DARPA have and will continue to play a lead role in developing enabling technologies. The CRDCs will interact with these agencies to be aware of new developments and to then help explore the implementation of the technologies.

However, because of resident expertise in some of the DOE labs, contributions to enabling technologies are possible. Examples include

- Free-electron laser photolithography
- Optical computing (electro-optical)
- Superconducting microelectronics

Federally-supported research in areas such as these would offer substantial leverage for U.S. industry in a collaborative R&D effort.

5. *General user program* -

(a) Provide user-friendly computing through networks to local communities for business, secondary schools, and universities.

(b) Demonstrate to the general user the value of user-friendly, distributed computing by highlighting pilot projects for members of the CRDC.

CRDC Selection and Funding

Because of the focus of S. 1976, we limit consideration of the CRDCs in this white paper to DOE laboratories as the lead organization in such Centers. However, in the distributed sense that has been outlined, the lead organization in a given CRDC would be linked to other principal participants such as universities, other labs, and industrial partners. It is possible that a CRDC would also prove to be effective if lead by an organization in another federal agency, such as NASA.

Within DOE, it would be prudent to choose the lead lab of a CRDC that has

- Current outstanding computing facilities
- Significant Grand Challenge applications
- History of working successfully with computing vendors

We see the possibility of 2-3 DOE Centers led by DOE labs. Although it is not the purview of S. 1976, we can also see the possibility of a total of 3-5 CRDCs nationally (including the DOE Centers). In the first year of the HPCP, only two CRDCs should be established in order to gain experience with the CRDC concept before expanding it.

Funding should leverage existing investments for hardware, software, networking, and applications. It is currently estimated that federal funding of about $30M per year (to cover activities of all participants combined in any specific Center) initially would allow the formation of a CRDC and permit significant contributions to be made to the national program in the areas outlined above. In the out years, annual funding would need to increase to about $40-50M per year for continued success. The estimated funding would provide for development and procurement of computing hardware and associated infrastructure, a graduate student fellowship program, operations, and research on both computing technology and solutions of Grand Challenge applications (and the integration of the two).

The federal funding provided to a CRDC would support the activities of the federal participants and the university participants. Industrial participants (both vendors and industrial users of computing) would cost share in the sense that they would pay for their own involvement - using the same model (but different funding requirements) used by the successful Superconductivity Pilot Centers that have been developed by the Department of Energy. However, a CRDC may choose to help support a vendor's development of some prototypical technology.

Response from Dr. John Connolly of the University of Kentucky to questions from the Committee on Energy and Natural Resources, as follow-up to a hearing on the "Department of Energy High Performance Computing Act of 1989" held on March 6, 1990.

Questions from Senator Ford:

Q: Please comment on whether there should be a lead agency or body on this overall effort, and if so, who that should be.

A: In order to meet the objectives of the High Performance Computing Initiative (HPCI), a coordinated program is necessary in the areas of hardware, software, networking and education. S. 1067 and S. 1976 recognize this need, but the organizational strategy outlined in the legislation isolates and fragments the initiative within different government agencies of the federal government. Such allocations are not necessarily ill advised, if adequate program integration can be ensured.

The FCCSET, which is presently designated as the coordinating organization, does not have the resources or the authority to integrate effectively the activities of the various federal agencies. Also, what coordination the FCCSET may be able to accomplish is limited to the federal sector. The participation of nongovernment organizations is a crucial ingredient for success. For example, for the national network, the participation of academia and the private sector are important if it is to meet the needs of its eventual users and complete its transition to commercialization.

The success of the HPCI will ultimately depend not on spending levels, but rather on the management structure placed upon it. This structure must integrate the components among federal agencies, and ensure the motivated participation of nongovernmental organizations. An integrating organization should be formed out of a formal alliance of government, academia and industry. This integrator should have connections to, and control over, all three sectors, but retain a degree of neutrality from them. As a first step in the HPCI, the OSTP should form this alliance.

Q: Does S.1976 place enough emphasis on education? If not, what should S.1976 emphasize?

A: S.1976 places a sufficient emphasis on the education of computational and computer scientists, but there are no provisions for training the other users of the network, such as those in industry, small business, K-12 students, etc. The level of computer literacy in the US is inadequate, which will become increasingly apparent when the national network is fully operational. Presently, efforts to increase computer literacy are fragmented and sporadic. The HPCI can provide a focus for such efforts.

Q: Should S.1976 require more or less access to Department of Energy supercomputers by Universities? Industry? Would the collaborative R&D Centers Dr. Hecker suggests give sufficient access?

A: The DoE now provides a small amount of access to universities for certain energy-specific projects, such as fusion research, particle physics and materials research. However, it is difficult for DoE to support projects, such as in the biomedical and environmental and other grand challenge areas, which are not connected with its mission. The collaborative R&D centers might be an appropriate mechanism to encourage coordination between the government, industry and academia, but their mission has to be broader than that of the DoE.

Q: Some of the witnesses testified that S.1976 could prevent flexibility by locking in specific expenditures within federal agencies. How should the bill be changed so as not to do this?

A: While the political expediency of placing the HPCI appropriations within each agency's annual budget is easily discernible, the flexibility of the integrating organization would be compromised. The means to avoid this inflexibility is to ensure that the authority for expenditures are vested in the integrating organization.

==========

Questions from Senator Domenici:

Q: In the evolution of TECHNET, the network I helped to start in New Mexico, the task of identifying database was critical to its success. Could you elaborate on what steps your respective institutions have taken in identifying users, other than academic institutions, for example, chambers of commerce, companies large and small, agencies from the private sector?

A: As an academic organization, the University of Kentucky is primarily focused on the educational aspects of high performance computing. However, we are about to initiate a new Advanced Science and Technology Commercialization Center (ASTeCC) which will include high performance computing and other areas, and which will concentrate on the technology transfer of knowledge to the sectors outside the university.

Also, the International Center for the Application of Information Technology (ICAIT) of Knoxville, Tennessee, which with we are associated, is currently preparing a broadbase opportunities and needs statement for the HPCI. This study will include input from non-academic organizations, and when it is published this summer, should answer your questions.

Q: Some of my colleagues who support the legislation in its present form, suggest that "hooking-up" institutions like yours will, itself, precipitate the educational process. This may be true for those like yourselves who already possess the necessary background and expertise, but could you tell me how others, who possess less background, will "interface" with the network? What would be needed for anyone owning a PC to be able to use this network?

A: Presently, a training and education component for network use is absent from the legislation. This omission should be corrected. Such a component would help encourage a broadbase of user support for the network as well as increase the computer literacy in the US. Input from non-technical users is vitally important in the early stages of the network design. Using the network should not be a complicated and onerous task, but rather a transparent user-friendly infrastructure which allows a PC-owner to access the vast array of resources. In establishing such an infrastructure, the legislation should reflect the network's potential impact on the ordinary user.

Q: One approach to a supercomputing network has been attempted by the National Science Foundation. It has 5 centers that constitute "nodes". Would you comment on the benefits and problems associated with this approach and contrast it to one which placed the hardware at one facility.

A: The NSFnet, which now has about 20 "backbone" nodes, which connects over 300 institutions around the US, and which now has a large number of international connections, has become the principal mechanism for information exchange and remote computer access for the global research community.

The national network should provide the infrastructure for information exchange among the entire community, which involves other sectors outside of research. If properly designed, the HPCI could shorten the time it takes to accomplish the transition of advanced technology into the mainstream.

In response to the question, the location of the hardware is eventually a moot question. The gigabaud network envisaged in this legislation would make the particular site irrelevant.

Q: Could you describe for me the present situation regarding the accessibility of the supercomputing centers at DOE facilities?

A: At present, DoE supercomputing facilities are not generally available to the academic community. There are three reasons for this--security, mission and cost. Most of the work done on DoE supercomputers is classified, so that the need to ensure security is significant. The DoE supercomputers are utilized for specialized

work connected to the DoE mission, and the demand for such specialized accessibility would be significant, and there is not a sufficient demand that it be done.

Q: In your opinion, what does this legislation need to properly address the question of network management? Who (or what) in your opinion is best suited for this?

A: The national network, as currently constituted, is inadequately managed. The network consists of a number of independent and badly coordinated interlocking subnetworks. No one is in charge and no one wants to be. An example of this lack of management is that any user who has access to the network has effectively unlimited access. This is very different from the control on supercomputer access, for which limited amounts of CPU cycles are assigned on the basis of merit. A similar access plan is required for the national network. This plan should be formulated and coordinated by a integrating organization which is independent of the federal agencies and which includes academic and industrial components.

The management of the network, and the rest of the HPCI, should include all interested parties. The integrating organization should be capable of coordinating tasks across the entire RDTE&C cycle (research, development, testing, evaluation and commercialization).

Appendix II

Additional Material Submitted for the Record

ON BEHALF OF THE

INSTITUTE OF ELECTRICAL AND ELECTRONICS ENGINEERS
UNITED STATES ACTIVITIES

ON THE

S.1976, THE DEPARTMENT OF ENERGY HIGH-PERFORMANCE
COMPUTING ACT OF 1989

BEFORE THE

UNITED STATES SENATE SUBCOMMITTEE ON ENERGY
RESEARCH AND DEVELOPMENT

OF THE

COMMITTEE ON ENERGY
AND NATURAL RESOURCES

STATEMENT ON BEHALF OF THE
INSTITUTE OF ELECTRICAL AND ELECTRONICS ENGINEERS
UNITED STATES ACTIVITIES
on
S.1976: THE DEPARTMENT OF ENERGY HIGH-PERFORMANCE COMPUTING ACT OF 1989
before the
United States Senate
Committee on Energy and Natural Resources
Subcommittee on Energy Research and Development

The IEEE is the world's largest engineering society representing over 315,000 worldwide, of whom approximately 250,000 live and work in the U.S. The IEEE seeks to advance the field of electrical, electronics, and computing engineering by disseminating scientific and technical information on a global basis. Within IEEE, the United States Activities Board addresses issues relating to the broad scope of communications services and develop recommendations on those issues. The IEEE-USA, Scientific Supercomputer Subcommittee is made up of individuals with a wide range of technical experience and expertise in all aspects of high performance computing. Appendix A is a list of the active participants on the sub committee and their affilations.

RECOMMENDATION:

IEEE strongly supports S.1976, the Department of Energy High-Performance Computing Act of 1989. All four of the major points of the legislation are vital to the success of the initiative: Section 4, on the National Research & Education Network; Section 5, on Software; Section 6, on High-Performance Computer Systems; and Section 7, on Basic Research and Education. All of these are important to the Department of Energy and its laboratories. We also support those activities requiring cooperation or coordination with other government agencies and the private sector, as addressed in Section 8, Government and Private Sector Cooperation.

The Department of Energy (DOE) has a long and strong history in the development and use of high performance computing. Its laboratories continue to support one of the largest and strongest group of sophisticated users working on a quite diverse set of important scientific and engineering problems. The DOE's ability to catalyze the growth of the high performance computing industry has been restricted during the last decade. This is due to diverse causes including technical and technological issues, the growth of the industry as a whole, external competition, infrastructures of government, industry, and universities not yet focussed on solving these problems; and legislation and government regulations in place to solve other problems.

Given the Department of Energy's potential to contribute to solving these problems, DOE can and should take on a strong leadership role in executing steps to correct the situation. S.1976 in its draft form appears to address these issues very well and its passage in the Senate should be pursued with diligence.

DISCUSSION:

Sensitive to the enormous importance of high performance computing to the economic well being of the United States and to its national security, the IEEE-USA Scientific Supercomputer Subcommittee has, over the last eight years, written a number papers striving to bring these issues to the attention of government, the supercomputing industry, and extant and potential industrial users of supercomputers.

The United States leadership in high performance computer technology is being threatened. The U.S. was the unchallenged leader in the development, manufacture and application of supercomputer technology until the mid

1980s. Today U.S. leadership in supercomputing is threatened by competition from outside the U.S. and more seriously by an inability to internally foster the healthy growth of an industry which by all objective accounts should play a major synergetic role with a large fraction of our industrial and research infrastructure. The problems are multi-faceted and include purely technical issues such as continuing development of hardware and software appropriate to high performance computing, and also include major issues of infrastructure, spanning government, industry, the universities, and the financial system.

In 1983, the IEEE made a study of the problems which were then beginning to develop. The report of that study urged the Federal government to take action to support U.S. leadership in supercomputing. Last year, in an effort to once again alert both the executive and legislative branches of government of our fears for the future of the supercomputer industry and more generally for the failure of the U.S. to more fully integrate high performance computing into its industrial and research activities, the IEEE subcommittee decided to update the 1983 study. The conclusion of that effort was that barring a few inconsequential details, the issues raised in the 1983 report were still valid. That conclusion and a summary of the important points and recommendations made was written into a Foreword dated May 1989 that accompanied a reissue of the 1983 report. Those conclusions and recommendations are still valid and are still very appropriate and quite pertinent to the thrust of S.1976. This Foreword is attached as Appendix B. (The full report is not included here, but is available from IEEE.)

Over the last half century, serial computing architecture, on which all current mature high performance computing systems are based, has been highly developed. Now, as we enter the last decade of the 20th century, we are in the process of introducing new parallel computer architectures in order to overcome limits on improved serial computer approaches imposed by technology and the speed of light. This leads to an exciting opportunity for innovative invention and the corollary that others may take advantage of a new paradigm where the historical leader may have no advantage.

As one surveys the field today, there are many companies striving to compete in the parallel computer architecture arena. Most start-up companies are not capable of maintaining a funding stream long enough to produce revenue from the production of high performance computers. The problems are serious because the development time and the associated costs for high performance computers are high. The investment typically required for developing new machines costs in the $100 million range and the typical price of each final manufactured product is in the $10 million range. Few of the competing companies in this marketplace have the resources to sustain such development costs. Start-ups requiring financing through venture capital or other financial institutions find it nearly impossible to find the resources to bring a product to marketplace. Furthermore, it is not just the processing units of a parallel system which are required, but also the associated high performance mass storage units, networking interfaces, the operating system and compiler support software, and last but not least the applications software. The current infrastructure of the producer and user communities is no longer working symbiotically as it has

in the past to bring computing systems to a useful operational state so that they are both effective for the user and revenue producing for the supplier.

DEPARTMENT OF ENERGY ROLE:

Since the early 1950's, the Department of Energy and its predecessor agencies, ERDA and the AEC, have been the premier users and the primary force in the high performance computing industry for the development of computers and computer systems. Over the years, the acquisitions of supercomputers by the Department of Energy laboratories and the articulation of their needs for future high performance computers in an open and ongoing dialogue with the designers and manufacturers of these systems gave the U.S. industry access to a highly competent customer and allowed the industry to develop. Although the designs of high performance computers were optimized for the needs of the Department of Energy laboratories, to a very high degree, high performance computers optimized for those problems serve equally well for a very wide range of scientific and engineering needs. It is, however, the example of a single set of applications which are well understood by a small cadre of experts willing and able to communicate their needs to the manufacturers and designers who have made this progress possible.

In the past, there has been a commitment on the part of the DOE laboratories in concert with the Department of Energy headquarters to purchase emerging equipment in advance of its production. This made for a very fine symbiotic relationship between the government laboratories and the computer manufacturers. Unfortunately, over the years, in trying to solve some of the problems of acquiring such complex and expensive systems, new

legislation as well as departmental and governmental guidelines have now inhibited these relationships. This has seriously undermined the synergistic mechanisms which over some years gave rise to the effective high performance computing industry in this country.

Although the Department of Energy laboratories have always developed software for applications appropriate to their missions, the new paradigms in computing, in particular parallel computing, make it appropriate to expend far greater resources in these developments in order to bring parallel processing approaches to an effective operational state. Here software institutes, focussed on a wide range of applications as well as on the system support software required to support these developments, are crucial to maintaining U.S. leadership. Although it is possible to establish such software institutes independently of existing organizations, the very best place to do this is within a vibrant organization which has, as its primary mission, a diverse set of problems which do require high performance computing. The Department of Energy laboratories represent a very good example of an appropriate infrastructure into which such software institutes should be established.

ATTACHMENTS

Appendix A (List of IEEE-USA Scientific Supercomputer Subcommittee Membership)

Appendix B (IEEE-USA Scientific Supercomputer Subcommittee Supercomputer Hardware Report Foreword, May 1989 [4 pages])

3-15-90 SCIENTIFIC SUPERCOMPUTER SUBCOMMITTEE ROSTER

Dr. Sidney Fernbach, Chairman
c/o Control Data Corporation
181 Metro Drive, Suite 700
San Jose, CA 95110
(408) 452-2140 (o)
(415) 837-9303 (h)
(408) 452-2151 (f)
EMail: fernback@lll-crg1inl.gov

Mr. Vito Bongiorno
Cray Research, Inc.
1333 Northland Drive
Mendota Heights, MN 55120
(612) 681-3620 (o)
(612) 681-3605 (f)
EMail: vb@ferrari.cray.com

Dr. Alfred E. Brenner
Supercomputing Research Center
17100 Science Drive
Bowie, MD 20715-4300
(301) 805-7400 (o)
(301) 805-7602 (f)
EMail: brenner@super.org

Dr. James F. Decker
DOE Office of Energy Research
Forrestal Building, MS 3F032
Washington, DC 20585
(202) 586-5430 (o)
(202) 586-4120 (f)

Mr. Dieter Fuss
Lawrence Livermore Nat'l Labs.
Mailstop L669
Box 808
Livermore, CA 94550
(415) 422-4027 (o)
(415) 423-4820 (f)
EMail: fuss@lll-crg.llul.gov

Mr. Mike Humphrey
Kendall Square Research
177 Bovet Road #600
San Mateo, CA 94402
(415) 341-2665, 345-7108 (o)
(415) 341-9709 (f)
EMail: mikehu@ksr.com

Prof. Duncan Lawrie
Center for Supercomputing R&D
University of Illinois
305 Talbot Laboratory
104 South Wright Street
Urbana, IL 61801
(217) 333-3373 (o)
(217) 244-1351 (f)
lawrie @ a.cs.uiuc.edu

Mr. Joel Maloff
CICNet, Inc.
535 W. William Street
Ann Arbor, MI 48103-4943
(313) 747-4272 (o)
(404) 330-5111 (voice mail)
EMail: maloff@merit.edu

Dr. Joanne L. Martin
IBM T.J. Watson Research Center
Post Office Box 218
Yorktown Hts., NY 10598
(914) 945-3285 - 945-3331 (o)
(914) 945-2141 (f)

Prof. Alan McAdams
Economic Strategy Inst.
1100 Connecticut Avenue, N.W.
Washington, DC 20036
(202) 728-0995 (o)
(202) 728-0998 (f)
EMail: amca@crnlgsm.edu

Dr. Norman Morse
Los Alamos National Laboratory
LANL CDO-MSB-200
Los Alamos, NM 87545
(505) 667-6164 (o)
(505) 667-1754 (f)

Dr. Kenneth W. Neves
Boeing Computer Services
ETA Div. M/S 7L-20
Post Office Box 24346
Seattle, WA 98124-0346
(206) 865-3504 (o)
(206) 865-2966 (f)
EMail: na.neves@.stanford.edu
kneves@atc.boeing.com

Dr. John M. Richardson
National Research Council
Room HA-370
2101 Constitution Avenue, NW
Washington, DC 20418
(202) 334-2844 (o)
(202) 334-2620, 334-2854 (f)
DialCom: 52:nrc2502
BITNET: JRICH @ nas
MCI Mail: John M. Richardson
NRC/SMP&ASP

Dr. John P. Riganati
Director for Systems Research
Supercomputing Research Center
17100 Science Drive
Bowie, MD 20715-4300
(301) 805-7450 (o)
(301) 977-6363 (h)
(301) 805-7602 (f)
EMail: riganati@super.org

Dr. Stewart Saphier
Asst. Prog. Manager for
Supercomputing Research
National Security Agency
9800 Savage Road
Ft. Meade, MD 20755-6000
ATTN: R5
(301) 859-6821 (o)

Dr. Paul B. Schneck
Director
Supercomputing Research Center
17100 Science Drive
Bowie, MD 20715-4300
(301) 805-7553 (o)
(301) 805-7602 (f)

Col. Will Stackhouse
Jet Propulsion Lab
4800 Oak Grove Drive
Mail Stop 79-23
Pasadena, CA 91109-8099
(818) 354-1668 (o) (FTS) 792-1668
(818) 354-0228 (f) (FTS) 792-0228

Mr. Lloyd Thorndyke
7409 Kellogg Avenue South
Edina, MN 55435
(612) 922-2901 (h)
(612) 921-6869 (f)

Dr. Paul A. Willis
Jet Propulsion Laboratory
MS 238-528
4800 Oak Grove Drive
Pasadena, CA 91109
(818) 354-1714 (o)
(818) 354-2825 (f)
COMP MAIL: P.WILLIS
ENERGY MGMT.SOCIETY

Dr. Robert Powers, CCIP Chairman
MCI Telecommunications Corp.
8283 Greensboro Drive
Mail Stop 1006/625
McLean, VA 22102
(703) 442-5295 (o)
(703) 734-2226 (f)
Pager: (202) 947-8131
Telex: 65-2728164MCIUW
MCI MAIL 272-8164
Internet: 2728164@MCIMail.COM

Ms. Deborah Rudolph, IEEE Staff
IEEE-USA
1828 L Steet, NW, Suite 1202
Washington, DC 20036
(202) 785-0017 (o)
(202) 785-0835 (f)

Dr. Arvid Larson, TPC Chairman
JJH, Inc.
2 Crystal Park, Suite 600
2121 Crystal Drive
Arlington, VA 22202
(703) 920-3435 (o)
(703) 893-4971 (h)
(703) 920-6758 (f)

SUPERCOMPUTER HARDWARE

A REPORT BY THE

IEEE-USA SCIENTIFIC SUPERCOMPUTER SUBCOMMITTEE

*NOVEMBER, 1983**

* *Updated in 1989 with new tables and a preface*

IEEE-USA SCIENTIFIC SUPERCOMPUTER SUBCOMMITTEE

SUPERCOMPUTER HARDWARE REPORT

FOREWORD

MAY, 1989

In 1983, the Scientific Supercomputer Committee of the IEEE United States Activities Board issued a report on the status of supercomputer development in the U.S. and abroad. It recommended several government actions that would support the U.S. leadership in this vital area. We reiterate these recommendations today. Recent Office of Science and Technology Policy (OSTP)/Federal Coordinating Committee for Science, Engineering, and Technology (FCCSET) reports urge similar actions.

In summary, the recommendations of the Committee in 1983 and in 1989 are that the Federal Government:

1. make a commitment to maintain leadership in scientific supercomputer development with a well defined, long range national program.

2. encourage both the continued development of scientific supercomputers and their availability to industry through the research community by adopting such measures as

   - direct support for development of new supercomputer generations,

   - direct support for establishing and operating several supercomputer centers for research, teaching and applications development,

   - tax credits or other financial incentives to supercomputer and related manufacturers,

   - tax credits or other financial incentives to supercomputer buyers and users,

   - increased direct support of related high risk research,

   - anti-trust relief to the functioning of industrial R&D consortia.

3. support of all technologies needed for the development of supercomputer systems.

4. designate a lead organization to coordinate the roles of the many Federal agencies which are critically dependent on supercomputer systems.

5. establish and extend national performance goals for the development of scientific supercomputer systems well into the 21st century. These long term goals should reflect the needs of Federally sponsored scientific research, both current and projected.

A readily implementable specific action that the Committee recommends can be taken now is the reinstatement of government procurement procedures which were very effective in nurturing the growth of the supercomputer industry during the early days of supercomputers. These procedures, often referred to as "buy before fly," allowed the federal procurement system to contract for supercomputers while still in design.

Since 1983, the following important changes have taken place in the supercomputer industry; 1) Control Data Corporation (CDC) spun off, subsequently reabsorbed, and on April 17, 1989, eliminated its supercomputer manufacturing facility, ETA Systems, Inc.; 2) a group of former Cray research employees left Cray to start a new company, Supercomputer Systems, Inc. (SSI), subsequently supported by IBM; 3) on May 15, 1989, Cray Research Corporation split off a separate entity, the Cray Computer Corporation, to manufacture the Cray-3 computer; 4) IBM has reentered the supercomputer business by the attachment of a vector facility to the 3090S mainframe. This system is in the low to middle performance range of supercomputer systems and had total sales at the end of 1988 of approximately 300 systems; 5) NEC announced on April 10, 1989, its new high performance line of computers including multi-processor models. This is the first announced Japanese product with more than a single processor.

With all these developments as background, the Committee thought its 1983 report ought to be reconsidered. Upon reviewing the 1983 report, the Committee finds that it remains both a valid and a current description of actions which must be undertaken by the U.S. in the supercomputer hardware area. We are reissuing the document, now under the auspices of the IEEE-USA Committee on Communications and Information Policy (as the Scientific Supercomputer Subcommittee), with an Addendum containing updated Tables and Figures on supercomputing. In 1983, our objective for these actions was to maintain our lead. In 1989, this objective remains but we are very close to having to <u>regain</u> rather than <u>maintain</u> our lead.

A positive move in the supercomputer arena since 1983 has been made by the National Science Foundation (NSF). As a result of a recommendation in a 1982 report by the Lax Committee, five supercomputer research centers were established by NSF and are now in operation. Unfortunately, it seems as if this number will not grow, and that even the five existing centers will be difficult to support in the manner deemed essential to keep them on the forefront of supercomputer developments.

Currently, there is an overwhelming momentum in this country to establish the high-end computer spectrum based on using numerous processors in parallel, each dealing with different portions of the same problem. The effective speed of the system is thereby presumed to be increased by n-fold where n is the number of processors. It is still true, however, that we have difficulty writing software to take advantage of such systems in an effective way. The IEEE Scientific Supercomputer Subcommittee is concerned about the inadequacy of software being provided by the U.S. manufacturers of supercomputers, and has expressed this concern in another report, which puts the software area into perspective and recommends actions.

The Japanese, our major competition in this arena, seem to be less interested in building parallel processing hardware. Their manufacturers endeavor to produce the fastest possible uniprocessors hoping to place them in parallel structures at a later time. While improvements in basic technology are not coming as fast as they used to, the Japanese are working in this area of development at a greater pace than we are.

A continuing need in supercomputers has been the growing demand for increased, directly-addressable memory. Semiconductor memory chips used in this country come mainly from Japan. The current price for the chips is much higher in the U.S. than in Japan, primarily because of Japan's leadership in the production of semiconductor memory components and the effect of U.S. trade agreements with Japan.

The market for supercomputers has grown significantly since 1983. At that time, 60 had been installed, mostly in the U.S. with approximately half at government sites. As of March 1989, approximately 400 (by Cray, ETA, Fujitsu, Hitachi, and NEC) have been installed and rapid growth continues, with the greatest rate of growth being in Japan using Japanese-manufactured systems.

Updated tables in the Addendum show the present status (Table 1) and planned activities (Table 2) in industrial supercomputer development. Table 3 addresses national efforts in Europe and Japan which have been aimed at development of supercomputer hardware and software. The revised table on national efforts presents the nature and scope of the current activities and the methods for putting results into practice. The United States has no comparable efforts. The graph shows the cumulative number of vector architecture supercomputers.

The OSTP/FCCSET report goes further in its recommendations than does the attached 1983 IEEE report. OSTP considers improvements in networking and long-term support for more basic research in areas affecting future development of supercomputers. This subcommittee wholeheartedly supports these additional activities, as well as the basic ideas described in the attached report.

###

The IEEE Scientific Supercomputer Subcommittee is a technology-policy subcommittee of the Committee on Communications and Information Policy of the Institute of Electrical and Electronics Engineers-United States Activities, 1111 19th Street, NW, Suite 608, Washington, DC 20036. The members of the Subcommittee are:

Sidney Fernbach, Chairman	Dieter Fuss	John P. Riganati
Lara Baker	Duncan H. Lawrie	Stewart Saphier
Vito Bongiorno	Alan McAdams	Paul B. Schneck
Alfred E. Brenner	Norman Morse	Lloyd Thorndyke
James F. Decker	Kenneth W. Neves	Hugh Walsh

John M. Richardson, Chairman, Committee on Communications and Information Policy

For more information, contact IEEE-USA at (202) 785-0017.

The IEEE-USA Scientific Supercomputer Subcommittee has prepared a number of other reports as follows: -

- SUPERCOMPUTER PERFORMANCE MEASUREMENT (1987)
- THE SIGNIFICANCE OF JAPANESE SUPERCOMPUTERS FOR U.S. TECHNOLOGICAL LEADERSHIP (1987)
- SOFTWARE FOR SUPERCOMPUTERS (1988)
- THE COMPUTER SPECTRUM: A PERSPECTIVE ON THE EVOLUTION OF COMPUTING (1988)
- U.S. SUPERCOMPUTER VULNERABILITY (1988)
- POSITION STATEMENTS WRITTEN BY THE SUBCOMMITTEE INCLUDE:

  -U.S. Supercomputer Vulnerability (1989)
  -Advanced Scientific Computing Capability Development (1983-Updated 1988)
  -The Semiconductor Trade Issue and Supercomputers (1987)
  -Semiconductor Trade Issues (1987)
  -In Support of the Report "A Research and Development Strategy for High Performance Computing" Issued by the Office of Science & Technology Policy (1988)

- BRIEFING POINTS ON SUPERCOMPUTERS
- SUPERCOMPUTING: AN INFORMAL GLOSSARY OF TERMS

For more information phone IEEE-USA at (202) 785-0017.

2111 Wilson Boulevard, Suite 531
Arlington, VA 22201
703-276-9068 FAX: 703-276-3580

Testimony by

Ralph E. Crafts, President

of the

Ada Software Alliance

on

S. 1976

before the

Subcommitee on Energy Research and Development

Committee on Energy and Natural Resources

U. S. Senate

Chairman Wendell H. Ford

March 6, 1990

Introduction

Mr. Chairman and Members of the Subcommittee, the Ada Software Alliance (ASA) is pleased and honored to submit this testimony regarding S. 1976, "Department of Energy High Performance Computing Act of 1989."

The Ada Software Alliance was formed in 1989 to promote and enhance the use of the Ada programming language and associated software engineering technology in any and all applications and programs where the use of Ada can benefit the consumer.

ASA strongly supports and applauds the goals and objectives of the High Performance Computing Plan (HPC). We believe that developing and implementing such a plan is critical to this country's long-term viability as a major competitor in global economic markets. This testimony is intended to offer perspective and suggestions as to how to maximize the effectiveness of the HPC legislation.

General Observations

The ASA wholeheartedly agrees with the bill's findings regarding the importance of high performance computing to this nation. The erosion of U.S. competitiveness in global computer markets is not due so much to a lack of innovation and creativity, but rather to a lack of engineering and application of critical

technologies. In other words, the U.S. holds a leadership position in creating advanced computer technology, but lags severely in applying it.

For example, in the electronics industry, loss of U.S. market share to Japanese companies has been the result of the Japanese applying American technology more quickly and efficiently than their U.S. counterparts. Typically, it takes 18-19 years for new technologies in this country to reach widespread usage in the marketplace. The lack of an engineering orientation, especially in the area of software technology, is one of the most debilitating factors in U.S. loss of market share.

The absence of an engineering focus contributes significantly to delays in the transfer of critical technologies to the "real world." As S. 1976 points out, government regulations and procurement practices act as enormous barriers to the rapid and efficient transfer of technology breakthroughs to actual applications and use. In the high performance computing arena, where the applicable technologies are very advanced, complex, and often abstract, the inability of key procurement agencies and personnel to understand and deal with these issues contributes significantly to the use of inadequate and obsolete regulations and directives.

The single common technology which applies to every facet of the HPC program is software. Software comprises the single most important enabling technology for all of the S. 1976 objectives. Conversely, software is also the most costly component in the development of any new computerized application. Fully 85 percent of the development costs of all computerized systems, including supercomputers, is expended on software. As noted earlier, the lack of an engineering orientation contributes greatly to the erosion of U.S. leadership in technology markets. The lack of engineering is more pronounced in the software arena than in any other segment addressed in S. 1976.

Software engineering is not the same as computer science. By comparison, computer science constitutes a small fraction of the disciplines, methods, and technologies encompassed by software engineering. Given the importance of software engineering to this country's ongoing success in the computer industry, it is very disturbing to note that there is not one single undergraduate degree program in software engineering in this entire nation. The insertion of software engineering into the HPC program agenda is essential to attaining the objectives delineated in S. 1976.

Suggestions and Recommendations

ASA believes that the HPC program can achieve its goals more effectively by including software engineering as a major component to be funded and improved. Software engineering provides the foundation for achieving reliable, reusable, and portable software for all computer applications. To realize the full potential of new computing architectures, such as massively parallel processing systems, software engineering is an absolutely essential ingredient for success. Thus, we believe there need to be specific references in S. 1976 to software engineering to highlight its significance.

Further, the participating agencies in the HPC program should focus their attention on the use of already existing standards in the software community. In particular, programming languages which are currently common standards between the U.S. military (MIL-STD), the American National Standards Institute (ANSI), and the International Standards Organization (ISO), should be emphasized. The legislation should emphasize this point.

The Subcommittee should be fully cognizant of the non-technical challenges associated with the HPC program. For example, the "NIH Syndrome" (Not Invented Here) is a very real phenomenon which could serve to delay or prevent progress. In

that regard, the Subcommittee should carefully consider where the HPC program is placed in the government hierarchy to minimize the perception that one particular agency has total control over the program and is realizing more benefits than other participating organizations. We are hopeful that the best features of S. 1976 will be meshed with S. 1067, Senator Gore's bill, to produce an overall governmental approach to the HPC initiative.

Conclusion

The Ada Software Alliance believes that the National High Performance Computing Program is long overdue and necessary to ensure America's position as a major competitor in global markets. ASA members also believe that the HPC could have a more significant impact by incorporating a focus on the actual application and engineering of technology in "real world" markets. To this end, use of existing common standards, particularly in software programming languages, would contribute significantly to the success of the HPC program.

Finally, the Alliance believes that the HPC needs to ensure awareness of the importance of the program at all levels of government, the Congress and the American people, to ensure the long-term viability and success of this critical initiative. The

ASA supports the HPC program, and stands ready to assist in its implementation.

Dr. Sidney Karin
Director
(619) 534-5075

SAN DIEGO SUPERCOMPUTER CENTER

GENERAL ATOMICS, P.O. BOX 85608, SAN DIEGO, CA 92138-5608

March 19, 1990
SK-90-011

Sen. Wendell H. Ford, Chairman
Subcommittee on Energy Research
and Development
Senate Committee on Energy and
Natural Resources
Washington, D.C. 20510

Dear Senator Ford:

Thank you for inviting me to testify before your Subcommittee on S. 1976, the DOE High Performance Computing Act of 1989. Unfortunately, as a result of a broken leg, I was unable to travel to Washington, D.C. for the hearing. However, Julie Shisler of my staff has been in contact with Paul Barnett regarding the submission of my testimony for the hearing record. To that end, enclosed are two copies of my prepared remarks as well as an attachment with specific comments about and suggestions for changes in S. 1976.

As the enclosed testimony notes, the ability of the United States to remain a leader in advancing science and technology requires a strong and sustained commitment to providing the best computational resources and widespread access to these resources. You are to be commended for taking a leadership role in assuring that a coordinated Federal plan is established and that the DOE has a well-focused role in the Federal High Performance Computing Program.

Again, that you again for the opportunity to provide my endorsement of S. 1976. If I can be of assistance to you in the future, please don't hesitate to call upon me.

Sincerely,

pc: Paul Barnett, Esq.
Senate Committee on Energy and
Natural Resources

LOCATED *at University of California, San Diego and administered by General Atomics*
CONSORTIUM MEMBERS: *The Agouron Institute • California Institute of Technology • National Optical Astronomy Observatories • Research Institute of Scripps Clinic • Salk Institute for Biological Studies • San Diego State University • Scripps Institution of Oceanography (UCSD) • Southwest Fisheries Center • Stanford University • University of California at Berkeley • University of California at Davis • University of California at Irvine • University of California at Los Angeles • University of California at Riverside • University of California at San Diego • University of California at San Francisco • University of California at Santa Barbara • University of California at Santa Cruz • University of Hawaii • University of Maryland • University of Michigan • University of Southern California • University of Utah • University of Washington • University of Wisconsin*

TESTIMONY OF
DR. SIDNEY KARIN
DIRECTOR
SAN DIEGO SUPERCOMPUTER CENTER

ON

S. 1976, THE DEPARTMENT OF ENERGY
HIGH-PERFORMANCE COMPUTING ACT OF 1989

BEFORE THE

SUBCOMMITTEE ON ENERGY RESEARCH AND DEVELOPMENT
OF THE
SENATE COMMITTEE ON ENERGY AND NATURAL RESOURCES

MARCH 6, 1990

TESTIMONY OF
DR. SIDNEY KARIN
DIRECTOR
SAN DIEGO SUPERCOMPUTER CENTER
ON
S. 1976, THE DEPARTMENT OF ENERGY
HIGH-PERFORMANCE COMPUTING ACT OF 1989
BEFORE THE
SUBCOMMITTEE ON ENERGY RESEARCH AND DEVELOPMENT
OF THE
SENATE COMMITTEE ON ENERGY AND NATURAL RESOURCES
MARCH 6, 1990

Mr. Chairman and Distinguished Senators:

My name is Sid Karin. I was asked to testify, but I couldn't get to the hearing because I broke my leg a week ago when I was trying to go very fast down a mountain. We skiers are taught early and reminded often to reassess our surroundings and respond to all the varying factors, not only to the snow conditions, but also to the terrain, the folks in front of you, and those "hot dogs" behind you! When we go really fast, we need to make very rapid adjustments or you lose control. The lesson I learned skiing is a lesson we're all learning in the field of high-performance computing. It was my fascination with going fast that drew me to high-performance computing in the first place, I suppose, but while I can break a leg, I'd rather my country didn't. So I thank you for the opportunity to add my two cents to the record of the testimony.

My title is director of the San Diego Supercomputer Center, one of the supercomputer centers established by the National Science Foundation in 1985. I have another title also, which is vice president of a San Diego research and development company called General Atomics, which runs the San Diego Supercomputer Center under the NSF Cooperative Agreement. I'm also an adjunct professor in the Computer Science and Engineering Department at the University of California, San Diego. Prior to donning these three hats, I had yet another title: manager of the User Service Center at General Atomics, that is, the User Service Center of the National Magnetic Fusion Energy Computer Center, one of the Department of Energy supercomputer centers, located at Livermore.

Our NSF supercomputer center, the one I direct today, was so quick to get off the ground five years ago precisely because it arose as a kind of spinoff of or evolution from the DOE supercomputer effort. It was the first NSF center to have experience with linking

users nationwide on a network, again because of its close relationship to the DOE effort. It's the product of a certain evolutionary path. I want to encourage your positive reading of S.1976 by drawing that evolutionary path, which is the path of high-performance computing, into the future. I have been where high-performance computing has been, and I think that gives me some insight into where it has to go.

The bill before you has four parts. The elements are establishment of a national network, development of improved software, support for high-performance computer systems, and basic research and education. This bill represents a balanced approach to continuing United States leadership in high-performance computing, and I am heartily in favor of funding all of these elements.

I should emphasize that I don't think we should be first in high-performance computing just because it would be nice to be first in as many areas as we possibly can. High-performance computing, as an enabling technology, is directly related to United States competitiveness in the world economy, and I want to explain how each part of this bill can be used to bolster and reinforce our competitive position.

The direct benefits of a national network are very obvious, as obvious to us today as the Interstate Highway System was to everyone in Eisenhower's time, when I was just a kid who loved fast cars.

Let me tell you how obvious it has become that we need a national network. When we started the San Diego center, we had before us the experience of the Magnetic Fusion Energy center, MFE, which was the first to link supercomputing researchers across the country by telephone and satellite on a true network. At all other nonclassified supercomputer centers, the remote users, and there weren't many, had to dial in and pay for a long-distance call.

When the San Diego center began, we knew we would need a network like MFE's, and we put it together that way, with links reaching from Maryland to Hawaii, with the result that we had more users working at our center between 1986 and early 1988, when the NSFnet backbone and regional networks began to take shape. We joined the NSFnet effort and began to phase out our MFEnet clone. Today NSFnet and its associated regional networks, like the California Education and Research Federation network that we run from San Diego, are providing service to increasing numbers of researchers around the country. I should add that high-performance computing is Darwinist: we've made the most of our inheritance, and we've adapted swiftly for survival.

So I think a larger national network linking all researchers at high speeds is an adaptation the United States requires for competitive survival, and this bill addresses it.

The funding required is not the largest part of the bill, but this initiative may well provide the most leverage.

If it has a defect, it is in its identification of the network as a "national research and education computer network" linking "government, industry, and the higher education community." Let's not make the mistake of isolating the research components of our economy from the engineering, design, and manufacturing sectors that benefit from the research. Let's link government, industry, higher education, and *production*—whether we call it manufacturing, trade, or commerce.

I'm suggesting that the various sectors of our economy can all evolve together. If they don't, we'll be in the same position as travelers were when the railroads began, having to hop out of trains at the border because the tracks in the next country were of a different gauge. Let's not keep the most productive sectors of our economy on a different gauge from research and education.

That leads me to another point about balance: high-performance computing is so important that it requires the cooperation of all agencies: not only DOE, but also NSF, NIH, NASA, the departments of Commerce and Agriculture, and others. I hope you will find a way to put in this bill or in another bill a mandate to all agencies to cooperate and to consider what the United States requires for competitive survival as they do so.

Since I'm supporting evolution here, let me point to the ultimate development. A computational superhighway should link everyone in the same way that telephones and fax machines do today. That idea is contained in the bill in the thought, here on page 8, that the government network should be phased out or turned over to the commercial networks when they can meet the needs. At that point, all of the elements of our society that *can* collaborate will be *able to* collaborate, and Japan Incorporated will be looking at U.S. Incorporated. As a number of editorialists have suggested in recent weeks, we may also need to face Germany, Incorporated, in the world market very soon. And in 1992, Europe, Incorporated.

That goal won't be met by networking alone. We need the other parts of this bill. The balance of the program is essential. Networks will let us talk to each other and compute together, but what will we compute, and how fast can we do it? And how can we assure that all of these elements work together to provide a robust computational infrastructure?

Let me bring in Darwin again to look at this. My center is a product of evolution. Its ancestor was a youthful DOE center that is still a vigorous citizen. Our center owes a lot to this parent, and to aunts and uncles on the DOE side as well. We obtained software, documentation, and much good advice from MFE. Our mass storage system, which we

have now commercialized under license from the DOE labs, was a software product from Los Alamos, with hardware from IBM. In addition, we have a similar technology transfer agreement with Lawrence Livermore for its UNIX-based file management system. As we have developed and interacted with our brother and sister NSF centers, we have been able to repay our parents and aunts and uncles. We have worked with the DOE labs on operating system issues. We are working with Los Alamos and Caltech on prototyping a gigabit per second network of the kind envisioned in this bill.

And we've also grown up and had kids of our own, and so have our brothers and sisters. We have fostered and participated in much important research, including research that speaks to one of the grandest challenges we face: maintaining American competitiveness. For example, in our center, we helped Gencorp to develop a new method for rapid injection molding of fiberglass and plastic auto body parts. The NSF center in Pittsburgh helped Alcoa to develop lighter aluminum cans, and the NSF center in Illinois helped Kodak develop new chemical processes for film developing. Moreover, we've had a vigorous program of research and technology transfer conducted jointly with members of the computer industry. At San Diego, this has included Cray, Digital Equipment, Supertek, Apple, and Alliant. This program has enabled these vendors to enhance their product offerings and improve their competitiveness in the world market.

We have taken very seriously our responsibility to educate researchers, students, and young people in the use of supercomputers and the opportunities of computational science. There's the Superkids program at MFE, the SuperQuest program at Cornell, the Kid Lab program at San Diego, and programs in Research Experiences for Undergraduates at all the NSF centers, in addition to fellowships and appointments for visiting scientists. All of these centers give workshops and hold meetings to spread technique, whether it's new math or new supercomputer graphics or new biochemistry. And we also clone ourselves. Key staff members of some of the new state-supported supercomputing centers have been drawn from the many experts trained by the DOE/DOD and NSF centers.

In other words, the NSF centers and the existing DOE centers have taken a leadership role in the efforts this bill seeks to support. Those efforts are the efforts that the present centers of supercomputing have been engaged in, multiplied thirty- or fortyfold. A layer of expertise, a computational mini-infrastructure, has already evolved. The efforts envisioned in this bill will have the most leverage if the funding is channeled through this layer. I hope the language referring to new centers will be looked at very closely when the time comes to implement this bill. A center is supposed to be central and is supposed to centralize, so it can act rapidly and take advantage of economies of scale. We should not repeat the

experience we already have; rather, we should build on it and very quickly evolve upward from it.

Darwinian Nature is profligate with mass extinctions and reformations of the plant and animal kingdoms, but, as Darwin himself pointed out, when humanity domesticates and breeds animals, it can choose to do better than Nature will do, and *faster* as well.

A Review of S. 1976
The Department of Energy
High Performance Computing Act of 1989

Section 2. (a) (5)

"In order to strengthen America's computer industry and to assist the entire manufacturing sector, the Federal Government must provide leadership in the development and application of high-performance computer technology."

We would encourage that the emphasis of this statement be on providing scientists and engineers in the manufacturing sector with the most advanced computational resources for translating results of research into reality, thereby hastening the development of new products which benefit us all. To support this concept the statement should therefore read "application of high performance computing technology".

"In particular, the Federal Government should create a National High-Performance Computing Program to support the development of a high capacity, national research and education computer network; facilitate the development of software for research, education, and industrial applications; continue to fund basic research; and provide for the training of computer scientists and computational scientists."

The final sentence in this section should also include the phrase "strengthen the existing federally-funded computational facilities."

Sec. 4. (b) (2)

"provide computer users at more than one thousand universities, Federal laboratories, and industrial laboratories with access to supercomputers, computer data bases, and other research facilities;"

This phrase should contain the word "facilities" to include the engineering, design, and manufacturing sectors, wherein the results of

computational research are applied to the development of new products and technologies. That is, after "industrial" strike "laboratories" and insert in lieu thereof "facilities".

Sec. 4. (b) (7)

"be phased out when commercial networks can meet the networking needs of American researchers."

The wording "be phased out" should be changed to "turned over to the private sector when . . .".

Finally, strengthen Section 4. (b) of the bill by adding the following:

"(8) foster the development and use of high-performance computing networks and systems by U.S. industry and universities."

In other words, this bill and those for other agencies should make provisions for a broader scope of responsibility in order to strengthen the national capability for high-performance computing capabilities, i.e., technology transfer, software development,

Sec. 4 (c) (2)

"provide for interconnection of existing computer networks run by the Department and other agencies, where appropriate."

Strike the phrase "where appropriate" and insert in lieu thereof ", as well as the regional networks." To assure the long-term success of the NREN, great emphasis should be placed on compatibility and reliability of all U.S. networks, including the Internet, i.e., the NSFnet Backbone, and mid-level regional networks.

Sec. 5. (a)

"the Secretary shall provide for research and development of high-performance computer software . . . [and] the development of

improved software tools and components to facilitate the development of software for high-performance computer systems."

The statement should be strengthened to include the following sentence: "Government/industry licensing agreements should be developed for federally-funded, unclassified software to be commercialized or otherwise made available to industry and universities."

Sec. 7. (a) (5)

"educate and train more reseachers in computer science and computational science by --

(A) making the national laboratories available to senior graduate students, postdoctoral fellows, and faculty from the Nation's universities;

(B) expanding summer science programs for high school students;

(D) establishing more cooperative research programs with the academic computational science community."

We whole-heartedly agree with the goals of A, B, and D, although we must note that these activities are the mission of the NSF Supercomputer Centers.

The NSF Centers support computational researchers across a broad array of scientific disciplines. In addition, the NSF Centers are all located on the campuses of major universities, therefore, the infrastructure and the experience to assist academic researchers with their education and training is in place and functioning well.

A, B, and D are worthy goals and could be enhanced by being implemented through a cooperative DOE/NSF effort, working through the existing NSF Supercomputer Centers.

Sec. 7. (a) (5)

(C) providing computer facilities to universities throughout the country:

The NSF Centers are currently providing broad access to advanced scientific computing facilites for the academic community not associated with the mission-oriented DOE, DOD and NASA facilities. Therefore, we believe that the DOE should address the basic research comunity, through the existing NSF Supercomputing Centers, chartered to meet that demand.

In addition, to assure the overall success of the U.S. high performance computing initiative, at some point in S. 1976, provisions should be added that will increase the United States' efforts to acquire new technologies that have been developed outside of this country. The bill should specifically include:

- Support for information gathering missions composed of representatives from the industrial, academic, and governmental communities to travel to other countries, such as Japan, to keep abreast of current high-performance computer research and technology.

- Support for translation and distribution of foreign language (in particular Japanese) journals and technical magazines.